z/OS Communications Server

IP Diagnosis Guide

Version 2 Release 2

GC27-3652-03

Note:
Before using this information and the product it supports, be sure to read the general information under "Notices" on page 1031.

This edition applies to Version 2 Release 2 of z/OS (5650-ZOS), and to subsequent releases and modifications until otherwise indicated in new editions.

IBM welcomes your comments. You can send us comments electronically by using one of the following methods:

Internet email:
comsvrcf@us.ibm.com

World Wide Web:
http://www.ibm.com/systems/z/os/zos/webqs.html

If you would like a reply, be sure to include your name, address, and telephone number. Make sure to include the following information 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 IBM Corporation 2000, 2015.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures . . . . . **xvii**

Tables . . . . . **xxi**

About this document . . . . . **xxiii**
Who should read this document xxiii
How this document is organized xxiii
How to use this document xxiii
Determining whether a publication is current xxiv
How to contact IBM service xxiv
Conventions and terminology that are used in this document xxiv
How to read a syntax diagram xxv
Prerequisite and related information xxviii

Summary of changes for IP Diagnosis Guide . . . . . **xxxiii**
Changes made in z/OS Version 2 Release 2 xxxiii
Changes made in z/OS Version 2 Release 1, as updated September 2014 xxxiv
Changes made in z/OS Version 2 Release 1, as updated December 2013 xxxiv
Summary of changes for z/OS Version 2 Release 1 xxxiv

Part 1. General diagnosis information **1**

Chapter 1. Overview of diagnosis procedure **3**
Steps for diagnosing problems 3

Chapter 2. Selecting tools and service aids **7**
How do I know which tool or service aid to select? 7
Selecting a dump 8
Selecting a trace 8
Selecting a service aid 12
Overview of available tools and service aids 13
Dumps 13
Traces 14
Display commands 16
System service aids 16
Methods for submitting documentation 17
Using AMATERSE 18
Using electronic transfer through email attachments 18
Necessary documentation. 18

Chapter 3. Diagnosing abends, loops, and hangs **21**
Analyzing abends 21
Analyzing loops 22
Steps for collecting documentation 22
Steps for analyzing hangs 23

Chapter 4. Diagnosing network connectivity problems **25**
Overview 25
Communicating through the correct stack 26
Steps for diagnosing problems connecting to a server 26
Steps for verifying server operation 27
Steps for verifying IP routing to a destination when not using policy-based routing 29
Steps for diagnosing problems with IP routing to a destination when using policy-based routing 31

© Copyright IBM Corp. 2000, 2015

Steps for verifying network interface operation . . . 33
Steps for verifying network access . . . 33
Tools for diagnosing network connectivity problems . . . 34
Using the Ping command . . . 34
Using the Netstat command . . . 39
Using the DISPLAY TCPIP,,OSAINFO command . . . 40
Using the Traceroute command . . . 41
Using SNMP remote Ping command . . . 41
Documentation for the IBM Support Center . . . 41

Part 2. Traces and control blocks . . . 43

Chapter 5. TCP/IP services traces and IPCS support . . . 45
Component trace . . . 45
Modifying options with the TRACE CT command . . . 46
IKE daemon trace (SYSTCPIK) . . . 58
Event trace (SYSTCPIP) for TCP/IP stacks and Telnet . . . 58
Specifying trace options . . . 59
Additional filters for SYSTCPIP . . . 67
Formatting event trace records for TCP/IP stacks and Telnet . . . 68
Socket API traces . . . 74
Packet trace (SYSTCPDA) for TCP/IP stacks . . . 93
The trace process . . . 94
Supported devices . . . 94
Starting packet trace . . . 95
Modifying options with VARY . . . 95
Formatting packet traces using IPCS . . . 96
Data trace (SYSTCPDA) for TCP/IP stacks . . . 146
Starting data trace . . . 146
Displaying data traces . . . 147
Formatting data traces using IPCS . . . 147
Intrusion Detection Services trace (SYSTCPIS) . . . 149
Restrictions . . . 150
CTRACE keywords on SYSTCPIS . . . 150
SYSTCPIS OPTIONS . . . 151
OPTIONS keywords . . . 152
IDS reports . . . 154
OSAENTA trace (SYSTCPOT) . . . 166
The OSAENTA trace process . . . 167
Starting OSAENTA trace . . . 168
Network security services (NSS) server trace (SYSTCPNS) . . . 170
Defense Manager daemon (DMD) trace (SYSTCPDM) . . . 170
OMPROUTE trace (SYSTCPRT) . . . 170
RESOLVER trace (SYSTCPRE) . . . 170
Configuration profile trace . . . 171

Chapter 6. IPCS subcommands for TCP/IP . . . 173
Types of subcommands . . . 173
TCPIPCS command . . . 176
Syntax . . . 177
Parameters . . . 177
Symbols defined . . . 179
TCPIPCS subcommands . . . 179
TCPIPCS API . . . 179
TCPIPCS CONFIG . . . 181
TCPIPCS CONNECTION . . . 182
TCPIPCS COUNTERS . . . 184
TCPIPCS DUAF . . . 186
TCPIPCS DUCB . . . 189
TCPIPCS FRCA . . . 192

TCPIPCS HASH . . . 194
TCPIPCS HEADER . . . 198
TCPIPCS HELP . . . 200
TCPIPCS IPSEC . . . 201
TCPIPCS LOCK . . . 204
TCPIPCS MAP . . . 205
TCPIPCS MTABLE . . . 208
TCPIPCS POLICY . . . 209
TCPIPCS PROFILE . . . 211
TCPIPCS PROTOCOL . . . 215
TCPIPCS RAW . . . 218
TCPIPCS ROUTE . . . 220
TCPIPCS SOCKET . . . 223
TCPIPCS STATE . . . 225
TCPIPCS STORAGE . . . 246
TCPIPCS STREAM . . . 248
TCPIPCS TCB . . . 250
TCPIPCS TELNET . . . 252
TCPIPCS TIMER . . . 254
TCPIPCS TRACE . . . 256
TCPIPCS TREE . . . 260
TCPIPCS TSDB . . . 263
TCPIPCS TSDX . . . 265
TCPIPCS TSEB . . . 266
TCPIPCS TTLS . . . 267
TCPIPCS UDP . . . 270
TCPIPCS VMCF . . . 273
TCPIPCS XCF . . . 274
ERRNO command . . . 276
Syntax . . . 276
Parameters . . . 276
Sample output of the ERRNO command . . . 277
ICMPHDR command . . . 278
Syntax . . . 278
Parameters . . . 279
Sample output of the ICMPHDR command . . . 279
IPHDR command . . . 280
Syntax . . . 280
Parameters . . . 280
Sample output of the IPHDR command . . . 281
RESOLVER command . . . 282
Syntax . . . 282
Parameters . . . 282
Sample output of the RESOLVER command . . . 283
SETPRINT command . . . 303
Syntax . . . 303
Parameters . . . 303
Sample output of the SETPRINT command . . . 303
SKMSG command . . . 303
Syntax . . . 303
Parameters . . . 304
Sample output of the SKMSG command . . . 304
TCPHDR command . . . 305
Syntax . . . 305
Parameters . . . 305
Sample output of the TCPHDR command . . . 306
TOD command . . . 306
Syntax . . . 307
Parameters . . . 307
Sample output of the TOD command . . . 307
UDPHDR command . . . 307

Syntax . . . 308
Parameters . . . 308
Sample output of the UDPHDR command . . . 308
Installing TCP/IP IPCS subcommands by using the panel interface . . . 309
Entering TCP/IP IPCS subcommands . . . 309
Steps for entering a TCP/IP IPCS subcommand . . . 310
Step for using the batch option . . . 311

Part 3. Diagnosing z/OS Communications Server components . . . 313

Chapter 7. Diagnosing problems with the z/OS Load Balancing Advisor . . . 315
Diagnostic data . . . 315
Diagnosing Advisor and Agent problems . . . 316
Abends . . . 316
Workload not distributed to a particular application . . . 316
Workload not distributed as expected . . . 317
Advisor or Agent appears to be hung . . . 318
Group names in displays are indecipherable . . . 319
Load balancer connection terminates unexpectedly . . . 319
Agent-Advisor connection terminates unexpectedly . . . 319
Automatic restart manager (ARM) registration failure . . . 320
Debug settings and corresponding syslogd priority levels . . . 320

Chapter 8. Diagnosing problems with the automated domain name registration application (ADNR) . . . 323
Diagnostic data . . . 323
Diagnosing ADNR problems . . . 325
Abends . . . 325
ADNR fails to initialize . . . 325
ADNR not communicating with the Global Workload Manager . . . 325
Automatic restart manager (ARM) registration failure . . . 327
ADNR not updating zones in a DNS server . . . 327
DNS name servers managed by ADNR contain incorrect or outdated data . . . 327
Diagnosing unresponsive zones . . . 328
ADNR appears to be hung . . . 330
ADNR connection to the GWM terminates unexpectedly . . . 330
Debug settings . . . 330

Chapter 9. Diagnosing IKE daemon problems . . . 333
Overview of diagnosing IKE daemon problems . . . 333
Diagnosing IKE daemon problems . . . 334
Initialization problems . . . 334
Problems establishing security associations . . . 334
Resource constraint problems . . . 349
Abends . . . 350
IKE daemon debug information . . . 350
Obtaining syslog debug information for the IKE daemon . . . 351
Obtaining debug information using PagentSyslogLevel . . . 351
TCP/IP services component trace for the IKE daemon . . . 351
Using CTRACE . . . 351
Formatting IKE daemon trace records . . . 355

Chapter 10. Diagnosing network security services (NSS) server problems . . . 357
Overview of diagnosing NSS server problems . . . 357
Common NSS server initialization problems . . . 357
NSS client connection problems . . . 358
NSS XMLAppliance client API return codes and reason codes . . . 360
Network security services server debug information . . . 367
Obtaining syslog debug information for the network security service server . . . 367

TCP/IP services component trace for the network security services (NSS) server 377
Using CTRACE. 378
Steps for enabling the CTRACE at network security service (NSS) server startup 379
Steps for enabling the CTRACE at network security services server startup. 379
Steps for disabling the CTRACE at network security services server startup 380
Step for enabling the CTRACE after the network security services server is started 380
Step for disabling the CTRACE after the network security services server is started 380
Step for displaying the CTRACE status. 380
Enabling CTRACE after network security services server initialization 381
Formatting network security services server trace records 381

Chapter 11. Diagnosing dynamic VIPA and sysplex problems 383
Overview of diagnosing sysplex distributor problems 383
Steps for diagnosing sysplex problems 384
Steps for diagnosing problems using DVIPAs in source IP address selection for TCPconnections problems . . . 396
Steps for diagnosing problems with the SYSPLEX-wide ephemeral port assignment for distributed DVIPAs . . . 397
Diagnosing problems with the SYSPLEX-wide ephemeral port assignment for EXPLICITBINDPORTRANGE processing 399
Steps for determining an optimal range for the EXPLICITBINDPORTRANGE parameter 400
Steps for diagnosing problems with EXPLICITBINDPORTRANGE processing 401
Diagnosing sysplex-wide security association (SWSA) problems 403
Steps for diagnosing sysplex-wide security association (SWSA) problems 403
Steps for diagnosing sysplex routing problems 409
Steps for diagnosing Tier 1 z/OS sysplex distribution problems 411
Steps for diagnosing Tier 1 non-z/OS sysplex distribution problems 413

Chapter 12. Diagnosing access control problems 417
Overview of access control support 418
Diagnosing multilevel security consistency check messages (EZD1215-EZD1234) 419
Steps for verifying the configuration. 419

Chapter 13. Diagnosing line print requester and daemon (LPR and LPD) problems 421
Diagnosing LPR client and LPD server problems 421
Abends 422
Steps for diagnosing timeouts, hangs, and waits. 422
Incorrect output 423
LPR client traces 425
Step for activating LPR client traces 425
Step for creating client trace output 425
LPD server traces 431
Step for activating server traces 431
Step for creating server trace output. 431

Chapter 14. Diagnosing File Transfer Protocol (FTP) problems 443
FTP server 443
Structural overview 443
Definitions and setup. 444
Error exit codes 444
Name considerations for z/OS UNIX FTP 444
Translation and data conversion support 446
DB2 query support 447
JES support 449
Logging FTP server activity 450
Common z/OS UNIX FTP problems. 451
Diagnosing FTP server problems with traces 464
Documenting server problems. 472
FTP client 473
Execution environments 474
Setup 474
Naming considerations 475

Directing the client to exit when an error occurs . . . 475
Translation and data conversion support . . . 475
File tagging support . . . 476
DB2 query support . . . 480
Restarting file transfers . . . 482
Diagnosing FTP connection and transfer failures with EZA2589E . . . 483
Diagnosing FTP client user exit problems . . . 489
Problems starting the client . . . 491
Problems logging in to the server . . . 492
Problems transferring data . . . 494
Other problems . . . 497
Diagnosing FTP client problems with tracing . . . 497
Where to find the FTP client trace . . . 498
Documenting FTP client problems . . . 499

Chapter 15. Diagnosing z/OS UNIX Telnet daemon (otelnetd) problems . . . 501
Common problems . . . 501
Debug traces . . . 502
Debug trace flows (netdata and ptydata) . . . 502
Debug trace examples (-t -D all) . . . 502
Cleaning up the utmp entries left from dead processes . . . 514

Chapter 16. Diagnosing Telnet problems . . . 515
General TN3270E Telnet server information . . . 515
TN3270E Telnet server definitions . . . 515
Diagnosing TN3270E Telnet server problems . . . 516
Abends (server) . . . 516
Logon problems (server) . . . 516
Session hangs (server) . . . 518
Incorrect output (server) . . . 520
Session outages (server) . . . 521
Special considerations when using SSL encryption support . . . 523
Telnet Component Trace data . . . 524
General Telnet client information . . . 524
Telnet client definitions . . . 524
Diagnosing Telnet client problems . . . 524
Abends (client) . . . 524
Session hangs (client) . . . 525
Incorrect output (client) . . . 527
Telnet client traces . . . 528
Step for starting Telnet client traces . . . 528
Trace example (client) . . . 528
Telnet commands and options . . . 532

Chapter 17. Diagnosing Simple Mail Transfer Protocol (SMTP) problems . . . 535
Sender SMTP . . . 535
Receiver SMTP . . . 535
SMTP environment . . . 535
SMTP definitions . . . 536
Diagnosing SMTP problems . . . 536
Abends . . . 536
Spooling problems . . . 537
SMTP does not deliver mail . . . 537
SMTP loop . . . 540
Mail item has incorrect output . . . 540
Forcing resolution of queued mail . . . 541
ADDRBLOK data set . . . 542
SMTP RESOLVER trace . . . 544

Chapter 18. Diagnosing z/OS UNIX sendmail and popper problems 547
Diagnostic aids for sendmail 547
Debugging switches 547
Additional diagnostic aids 549
Diagnostic aids for IPv6 support 551
Diagnostic aids for AT-TLS support 552
Diagnostic aids for mail filter support 552
Hints and troubleshooting sendmail message submission program (MSP) file submit.cf 553
Diagnostic aids for popper 554

Chapter 19. Diagnosing SNALINK LU0 problems 557
Definitions 557
Problem diagnosis 557
Abends 558
Session hangs 558
Session outages 560
Traces 561
Using IP packet trace 561
SNALINK LU0 DEBUG trace 561

Chapter 20. Diagnosing SNALINK LU6.2 problems 565
Steps for setting up a SNALINK LU6.2 network 566
Common configuration mistakes 567
Diagnosing problems 568
Quick checklist for common problems 568
Problems starting the SNALINK LU6.2 address space 569
DLC connection problems 571
Network connection establishment problems 574
Network connection loss problems 576
Data loss problems 577
Data corruption problems 579
Documentation references for problem diagnosis 579
Using NETSTAT 580
Using the SNALINK LU6.2 subcommand 580
Useful VTAM operations 581
Traces 584
Using SNALINK LU6.2 internal traces 584
Using IP packet trace 587
TCP/IP internal traces 588
VTAM buffer traces 589
Finding abend and sense code documentation 589
Finding error message documentation 589

Chapter 21. Diagnosing REXEC, REXECD, and RSH problems 591
General information about REXEC and RSH 591
Documentation for REXEC problem diagnosis 591
TSO console log 592
Activating the REXEC debug trace 592
REXEC trace example and explanation 592
RSH trace example and explanation 592
General information about REXECD 592
Documentation for REXECD problem diagnosis 593
MVS system console log 593
Starting REXECD server traces 593
Example of an REXECD trace of a client using the SEND command 593
Example trace of an RSH client using the SEND command 595
Example trace to the JES spool file of the server 596

Chapter 22. Diagnosing z/OS UNIX REXEC, RSH, REXECD, and RSHD problems . . . 599
Setting up the inetd configuration file 599

Diagnosing z/OS UNIX REXEC 600
Activating the z/OS UNIX REXEC debug trace 600
z/OS UNIX REXEC trace example and explanation 600
Diagnosing z/OS UNIX RSH 601
Step for activating the z/OS UNIX RSH debug trace 601
Step for invoking z/OS UNIX RSH trace 601
Diagnosing z/OS UNIX REXECD 602
Activating the z/OS UNIX REXECD debug trace 602
z/OS UNIX REXECD trace example and explanation 602
Diagnosing z/OS UNIX RSHD 602
Step for activating the z/OS UNIX RSHD debug trace. 603
z/OS UNIX RSHD trace example and explanation 603

Chapter 23. Diagnosing X Window System and Motif problems 605
Trace output when XWTRACE=2. 605
Trace output when XWTRACELC=2. 606

Chapter 24. Diagnosing Simple Network Management Protocol (SNMP) problems . . . 609
Overview. 609
Management information base (MIB) 609
PDUs 609
Functional components 610
Definitions 611
z/OS UNIX **snmp** command 611
SNMP agent. 611
TCP/IP subagent 612
OMPROUTE subagent 612
Network SLAPM2 subagent 612
TN3270E Telnet subagent 613
SNMP socket call settings 613
Trap forwarder daemon 613
Diagnosing SNMP problems 613
Abends 614
SNMP connection problems 614
Incorrect output 621
No response from the SNMP agent 627
SNMPv3 AES encryption or decryption failures 629
Report received from SNMP agent 630
0.0.0.0 address in traps from the SNMP agent 631
I/O error for SNMP PING 631
Traps not forwarded by trap forwarder daemon. 632
Incorrect address in forwarded trap 632
SNMP traces 633
Starting manager traces 633
Starting SNMP agent traces. 634
Starting TCP/IP subagent traces 635
Starting OMPROUTE subagent traces 635
Starting Network SLAPM2 subagent traces 635
Starting TN3270E Telnet subagent traces 636
Starting TRAPFWD traces 636
Trace examples and explanations 637

Chapter 25. Diagnosing Policy Agent problems 657
Overview. 657
QoS policy 658
QoS policy scope 658
Import services. 659
Gathering diagnostic information. 659
Diagnosing Policy Agent problems 661
Initialization problems 662

Policy definition problems . . . 663
Policy client connection problems . . . 668
Policy client retrieval problems . . . 671
Import requestor connection problems . . . 673
Import requestor retrieval problems . . . 678
LDAP object retrieval problems . . . 678
LDAP object storage problems. . . . 680
Policy Agent and Sysplex distribution problems . . . 681
Memory allocation/leakage problems . . . 682

Chapter 26. Diagnosing RSVP agent problems . . . 685
Overview. . . . 685
Reservation types, styles, and objects . . . 686
Policies and RSVP processing . . . 687
Gathering diagnostic information. . . . 688
Diagnosing RSVP agent problems . . . 688
Initialization problems . . . 688
Application problems. . . . 689
Policy problems . . . 689
Example log file . . . 690

Chapter 27. Diagnosing intrusion detection problems . . . 699
Overview. . . . 699
Diagnosing IDS policy problems . . . 699
Step for determining which IDS policies are active in Policy Agent . . . 699
Step for determining how your IDS policies have been mapped by the stack . . . 700
Diagnosing IDS output problems . . . 700
Steps for determining why IDS syslogd output is missing . . . 700
IDS console output . . . 700
IDS packet trace output . . . 701
Unusual conditions . . . 701
Diagnosing TRMD problems . . . 704
Documentation for the IBM Software Support Center . . . 704

Chapter 28. Diagnosing Application Transparent Transport Layer Security (AT-TLS) 707
Common AT-TLS startup errors . . . 707
Steps for diagnosing AT-TLS problems . . . 707
AT-TLS traces . . . 709
Sample AT-TLS trace . . . 710
AT-TLS return codes . . . 712
SIOCTTLSCTL ioctl return codes . . . 718

Chapter 29. Diagnosing Defense Manager daemon problems . . . 721
Overview of diagnosing Defense Manager daemon problems . . . 721
Defense Manager daemon debug information . . . 727
Abends during DMD processing . . . 727
DMD error codes . . . 727
TCP/IP services component trace for the Defense Manager daemon . . . 728
CTRACE options . . . 728
Enabling CTRACE at Defense Manager daemon startup . . . 729
Steps for enabling the CTRACE at Defense Manager daemon startup. . . . 731
Steps for disabling the CTRACE at Defense Manager daemon startup . . . 731
Step for enabling the CTRACE after the Defense Manager daemon has started . . . 731
Step for disabling the CTRACE after the Defense Manager daemon is started . . . 731
Displaying the CTRACE status . . . 732
Enabling CTRACE after Defense Manager daemon initialization . . . 732
Formatting Defense Manager daemon trace records . . . 732

Chapter 30. Diagnosing IP security and defensive filter problems . . . 733
Overview of diagnosing IP security and defensive filter problems . . . 733

Steps for diagnosing IP security problems . . . 734
Steps for diagnosing defensive filter problems . . . 735
Steps for diagnosing the cause for missing IP security or defensive filter syslogd output . . . 737
Steps for verifying IP security and defensive filter operation . . . 741
Steps for verifying manual IPSec protection . . . 744
Steps for verifying dynamic IPSec protection . . . 746
Steps for verifying IP security policy or defensive filter enforcement . . . 749
Steps for verifying IPSec processing on zIIP . . . 753
Determining the Workload Manager service class associated with IPSec workload being processed on zIIP . . . 753
Tools for diagnosing IP security and defensive filter problems . . . 753
Using the **ipsec** command . . . 753
Using the **pasearch** command . . . 757
Using syslog messages . . . 757

Chapter 31. Diagnosing OMPROUTE problems . . . 759
Overview . . . 759
Definitions . . . 762
Diagnosing OMPROUTE problems . . . 762
Abends . . . 762
OMPROUTE connection problems . . . 762
Routing failures . . . 762
Adjacency failures . . . 764
Malformed and unexpected packets . . . 773
OMPROUTE traces and debug information . . . 775
Starting OMPROUTE tracing and debugging from the z/OS UNIX System Services shell . . . 775
Starting OMPROUTE tracing and debugging from an MVS cataloged procedure or AUTOLOG . . . 776
Starting OMPROUTE tracing and debugging using the MODIFY command . . . 776
Destination of OMPROUTE trace and debug output . . . 777
Sample OMPROUTE trace output . . . 777
TCP/IP services component trace for OMPROUTE . . . 786
Specifying trace options . . . 787
Formatting OMPROUTE trace records . . . 789
Commands to enable, disable, and display the status of the OMPROUTE CTRACE . . . 790
Steps for enabling the CTRACE at OMPROUTE startup . . . 790
Steps for disabling the CTRACE at OMPROUTE startup . . . 790
Steps for enabling the CTRACE after OMPROUTE is started . . . 790
Steps for disabling the CTRACE after OMPROUTE is started . . . 791
Step for displaying the CTRACE status . . . 791

Chapter 32. Diagnosing NCPROUTE problems . . . 793
Definitions . . . 796
Diagnosing NCPROUTE problems . . . 797
Abends . . . 797
Connection problems . . . 797
Analyzing routing failures . . . 801
Incorrect output . . . 803
Session outages . . . 804
NCPROUTE traces . . . 805
Activating NCPROUTE global traces . . . 806
Activating NCPROUTE selective traces . . . 806
NCPROUTE trace example and explanation . . . 807

Chapter 33. Diagnosing X.25 NPSI problems . . . 817
Operation . . . 818
Configuration requirements . . . 819
RACF/Security Manager requirement . . . 819
VTAM considerations . . . 819
NPSI considerations . . . 819
Sources of diagnostic information . . . 820
X.25 trace examples . . . 820

Normal incoming call, TRACE OFF . . . 821
Normal incoming call, TRACE DATA . . . 821
Normal outgoing call, TRACE CONTROL . . . 822
Results of LIST command . . . 823
Termination by TCPIP STOP device . . . 823
Steps for diagnosing logon problems . . . 823
Session hangs . . . 824
Helpful hints . . . 825
Documentation requirements . . . 825

Chapter 34. Diagnosing IMS problems . . . 827
Steps for setting up the IMS TCP/IP services socket interface system . . . 829
Common configuration mistakes . . . 830
Quick checklist for common problems . . . 831
Component problems . . . 831
Connection problems . . . 833
Error message and return code problems . . . 837
Socket data protocol problems . . . 838
IMS transaction build problems . . . 841
IMS database problems . . . 842
Documentation references for problem diagnosis . . . 844
Traces . . . 845
Using NETSTAT . . . 846
Where to find return code documentation . . . 847
Where to find error message documentation . . . 847

Chapter 35. Diagnosing VMCF/TNF/IUCV problems . . . 849
Diagnosing restartable VMCF/TNF problems . . . 849
VMCF or TNF fail to initialize . . . 849
Abends 0D5 and 0D6 . . . 849
Steps for diagnosing no response to commands . . . 849
VMCF or TNF does not stop . . . 850
Diagnosing VMCF/IUCV problems with the TSO MVPXDISP command . . . 850

Chapter 36. Diagnosing problems with IP CICS sockets . . . 855
Diagnostic data . . . 855
Initialization problems . . . 856
Steps for diagnosing CICS socket interface not initialized . . . 856
Steps for diagnosing CICS listener not initialized . . . 856
No CICS sockets messages issued . . . 857
Steps for diagnosing TCP/IP clients unable to connect . . . 857
Steps for diagnosing child-server transactions not starting . . . 858
CICS sockets application problems . . . 858
Steps for diagnosing hung CICS tasks . . . 858
Hung CICS region . . . 858
Errors on socket calls . . . 859
CICS shutdown hangs . . . 859
CICS sockets control blocks . . . 859
Task interface element . . . 859
Global work area . . . 859
CICS trace . . . 860
Steps for displaying the internal trace . . . 860

Chapter 37. Diagnosing problems with SMC-R . . . 861
Switch configuration issues . . . 861
Physical network ID configuration issues . . . 862
No associated subnet mask . . . 862
PFID status remains STARTING . . . 863
Problem with SMC-R interaction with security function . . . 863

Chapter 38. Diagnosing problems with Express Logon 865
Analyzing start problems with the DCAS 866
Analyzing client interface problems 867

Chapter 39. Diagnosing resolver problems 869
Steps for resolving the hostname 869
Steps for resolving caching problems 872
Steps for responding to message EZZ9308E 875
Steps for responding to message EZZ9311E 879
TRACE RESOLVER 884
Specifying the Trace Resolver output location 887
Interpreting the Trace Resolver output 889
CTRACE - RESOLVER 907
CTRACE TRACERES 909
Enabling collection of CTRACE TRACERES records 910
Disabling collection of CTRACE TRACERES records 910
Formatting the CTRACE TRACERES option 911

Chapter 40. Diagnosing Simple Network Time Protocol (SNTP) problems 913
Activating the SNTPD debug trace 913
Abends 913
Steps for stopping SNTPD 913
Sample SNTPD debug output 914

Chapter 41. Diagnosing Communications Server SMTP application problems 915
Gathering diagnostic information 915
Steps for gathering log information 915
Resolving initialization or logging problems 916
Resolving SMTPNOTE CLIST problems 917
Diagnosing and resolving Resolver problems 918
Resolving problems from the JES spool data set 919
Resolving mail problems 923
Target server problems 923
Mail problems 925
Verifying mail workload compatibility with SMTPD 932
Resolving MODIFY command problems 933
Diagnosing checkpoint problems 933
Monitoring resources used 934
Monitoring target servers 938
ExtendedRetry Mail Directory 940

Chapter 42. Diagnosing storage abends and storage growth 941
Storage definitions. 941
Monitoring storage utilization 942
Limiting TCP/IP common and private storage utilization 943
Limiting CSM storage utilization 943
Storage messages 944
Sysplex Problem Detection and Recovery (SPDR) storage messages 945
Abends 945
Diagnosing a storage problem 946
Steps for reviewing a storage problem 946
Steps for reviewing data on the TCP queues 947
Steps for reviewing data on the queue if you are using IDS 949
Steps for reviewing data on the queue if you are not using IDS 950
Collecting documentation to submit to the IBM Support Center 951

Part 4. Appendixes 953

Appendix A. Overview of internetworking . . . 955
Maximum transmission unit (MTU) . . . 956
Fiber Distributed Data Interface (FDDI) . . . 957
Token-Ring IEEE 802.5 . . . 958
IEEE 802.3 . . . 959
Ethernet — DIX V2 . . . 959
Subnetwork Access Protocol (SNAP) . . . 960
IP routing . . . 961
Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) . . . 961
Internet Protocol Version 4 (IPv4) . . . 962
Internet Protocol Version 6 (IPv6) . . . 964
Direct routing . . . 964
Indirect routing . . . 965
Simplified IP datagram routing algorithm . . . 965
IPv4 subnetting . . . 966
IPv6 prefixes . . . 967
Simplified IP datagram routing algorithm with subnets . . . 967
Static routing . . . 969
Dynamic routing . . . 969
IPv4 . . . 970
IPv6 . . . 970

Appendix B. IKE protocol details . . . 971
IKE version 1 protocol . . . 971
Overview of negotiating IKEv1 security associations . . . 971
Phase 1 . . . 971
Phase 2 . . . 977
ISAKMP Main mode limitations . . . 982
IKE version 2 protocol . . . 983
Overview of negotiating IKEv2 security associations . . . 983
Initial exchanges . . . 984
Child SA activation . . . 988
Key exchange limitations . . . 990
Traversing a NAT . . . 992

Appendix C. IBM Health Checker for z/OS . . . 999

Appendix D. Related protocol specifications . . . 1003

Appendix E. Accessibility . . . 1027

Notices . . . 1031
Programming interface information . . . 1039
Policy for unsupported hardware . . . 1039
Trademarks . . . 1039

Bibliography . . . 1041

Index . . . 1045

Communicating your comments to IBM . . . 1057

Figures

1. Overview of the diagnosis procedure . . . 4
2. Example of output from the IPCS SYSTRACE command . . . 23
3. Overview of verifying server operation . . . 28
4. Overview of verifying IP routing to a destination . . . 30
5. Overview of verifying network interface operation . . . 33
6. Overview of verifying network access. . . . 34
7. Example of DUMP command for TCP/IP stack . . . 50
8. Example of DUMP command for OMPROUTE. . . . 50
9. Example of DUMP command for RESOLVER . . . 51
10. Example of DUMP command for TELNET . . . 51
11. IPCS CTRACE . . . 54
12. SYS1.PARMLIB member CTIEZB00 . . . 61
13. Start of component trace full format . . . 71
14. Component trace full format showing character interpretation of fields . . . 72
15. TCP/IP networking API relationship on z/OS . . . 75
16. Data trace record. . . . 93
17. SOCKAPI trace record.. . . . 93
18. Control and data flow in the IP packet tracing facility . . . 94
19. Example of a SUMMARY report . . . 121
20. Format report example . . . 129
21. Data trace: Single entry . . . 147
22. Data trace with multiple entries . . . 147
23. FORMAT option output . . . 148
24. SYS1.PARMLIB member CTIIDS00 . . . 150
25. Control and data flow in the OSAENTA tracing facility . . . 167
26. Main menu for TCP/IP IPCS subcommands. . . . 310
27. IPCS primary option menu . . . 310
28. IPCS subcommand entry panel with a TCP/IP IPCS subcommand entered . . . 311
29. NAT keep alive message . . . 348
30. SYS1.PARMLIB member CTIIKE00 . . . 354
31. Netstat VIPADCFG/-F example . . . 386
32. Netstat CONFIG/-f example . . . 387
33. D WLM,SYSTEMS example . . . 388
34. Netstat VIPADYN/-v example . . . 388
35. Netstat VIPADyn/-v example . . . 389
36. Sysplex VIPADyn example . . . 390
37. Netstat VDPT/-O example . . . 393
38. d tcpip,tcpcs,net,vcrt,detail example . . . 395
39. Netstat ALLConn/-a example with IPAddr/-I filter value of 201.2.1.12 . . . 395
40. Netstat CONFIG/-f example . . . 396
41. Netstat SRCIP/-J example . . . 397
42. Diagnosing SYSPLEXPORTS problems . . . 398
43. VTAM NET,STATS example. . . . 399
44. VTAM NET,STATS,LIST=0 example . . . 400
45. ipsec -f example . . . 404
46. ipsec -y example . . . 405
47. ipsec -y display -s example . . . 407
48. Netstat VIPADyn example . . . 409
49. Netstat VCRT/-V detail example . . . 411
50. Example of LPR trace output (Part 1 of 2) . . . 426
51. Example of LPR trace output (Part 2 of 2) . . . 427
52. Example of LPR trace with filter x option . . . 428
53. Example of LPR output with unknown printer . . . 429
54. Example of LPR trace with JNUM, LANDSCAPE, and TRACE options . . . 430
55. Example of LPR trace with XLATE option . . . 431

© Copyright IBM Corp. 2000, 2015

56. Example of LPD trace specified with the DEBUG option . . . 432
57. Example of an LPD server trace of a failing job . . . 437
58. Example of an LPD server trace for a remote print request . . . 439
59. Trace between the Telnet client, parent, and child . . . 502
60. z/OS UNIX Telnet trace using -t -D all . . . 503
61. Telnet client trace . . . 529
62. SMTP environment . . . 536
63. Example of RESOLVER trace output . . . 545
64. Example of a SNALINK LU0 DEBUG trace . . . 562
65. Components of a SNALINK LU6.2 connection on MVS . . . 565
66. Sample MVS System Console Messages on SNALINK LU6.2 Address Space Startup . . . 567
67. NETSTAT DEVLINKS output example . . . 580
68. LIST MODIFY subcommand output example . . . 581
69. DISPLAY subcommand output example for connectable LU . . . 583
70. DISPLAY subcommand output example for active LU . . . 583
71. SNALINK LU6.2 internal trace output . . . 585
72. A CTRACE formatted packet trace record . . . 588
73. Remote execution protocol principle . . . 591
74. Example of an REXEC trace. . . . 592
75. Example of an RSH trace . . . 592
76. Example of an REXECD trace of a client using a SEND command . . . 594
77. Example of a trace of an RSH client using a SEND command . . . 595
78. Example trace to the JES spool file of the server . . . 597
79. Adding applications to /etc/inetd.conf . . . 599
80. Setting traces in /etc/inetd.conf . . . 600
81. Example of X Application trace output when XWTRACE=2 . . . 606
82. Example of X Application trace output when XWTRACELC=2 . . . 607
83. SNMP agent response trace . . . 638
84. SNMP agent trace of unsuccessful initialization . . . 638
85. SNMP messages and agent trace for nonmatching key . . . 638
86. SNMP messages and agent trace when data not in defined view . . . 639
87. SNMP subagent trace . . . 640
88. SNMP query engine traces . . . 642
89. SNMP IUCV communication traces . . . 652
90. TRAPFWD trace . . . 656
91. RSVP Agent processing log . . . 691
92. Example trace of a generic server processing . . . 711
93. SYS1.PARMLIB member CTIDMD00 . . . 730
94. Overview of verifying IP security operation . . . 742
95. Overview of verifying manual IPSec protection . . . 745
96. Overview of verifying dynamic IPSec protection . . . 747
97. Overview of verifying IP security policy enforcement . . . 750
98. Sample OMPROUTE Trace Output . . . 778
99. Sample IPv6 OMPROUTE Trace Output . . . 786
100. SYS1.PARMLIB member CTIORA00 . . . 788
101. NCPROUTE environment . . . 793
102. NCPROUTE trace . . . 808
103. X.25 NPSI environment . . . 818
104. Components of the IMS TCP/IP services socket interface system . . . 828
105. MVPXDISP sample output using the *userid* parameter . . . 851
106. MVPXDISP sample output using the ISAQ parameter . . . 853
107. Error report example . . . 922
108. Example of undeliverable mail notification. . . . 926
109. Routers and bridges within an internet . . . 956
110. Relationship of MTU to frame size . . . 957
111. Format of an IEEE 802.5 token-ring frame . . . 958
112. Format of an IEEE 802.3 frame. . . . 959
113. Format of an Ethernet V2 frame . . . 960
114. SNAP header . . . 960
115. Routing and bridging . . . 962
116. Classes of IPv4 addresses . . . 963

117. Determining the class of an IPv4 address 963
118. General IP routing algorithm 965
119. Subnetting scheme 966
120. Routing algorithm with subnets 968
121. Example of resolving a subnet route 969
122. Main mode exchange 972
123. Aggressive mode exchange 973
124. Interpreting phase 1 SA states in Main mode 975
125. Interpreting phase 1 SA states in Aggressive mode 976
126. Quick mode exchange messages 977
127. Quick mode exchange with commit-bit support 979
128. Quick (phase 2) SA states without commit-bit support 980
129. Quick (phase 2) SA states with commit-bit support 981
130. IKEv2 initial exchanges 984
131. Interpreting IKEv2 IKE SA states 987
132. IKEv2 CREATE_CHILD_SA exchange 988
133. IKEv2 Child SA states. 990
134. IKE packet with and without the non-ESP marker 997

Tables

1. Selecting a service aid . . . 7
2. Selecting a dump . . . 8
3. Selecting a trace . . . 8
4. Description of dumps . . . 13
5. Description of traces . . . 14
6. Description of service aids . . . 16
7. TCP/IP component name and release level . . . 19
8. Types of abends . . . 21
9. Diagnostic commands . . . 26
10. Diagnosis of a timeout . . . 38
11. Trace options and groups . . . 61
12. Data types . . . 67
13. IP address and port filtering effect on different types of socket API calls . . . 79
14. Flags that apply to IP or SNA packets . . . 109
15. TCP/IP IPCS commands . . . 173
16. Available debug levels and associated syslogd priority levels . . . 320
17. Establishing security associations problems . . . 335
18. Common problems when IKED, running as an NSS client, is unable to obtain services from the NSS server 343
19. Resource constraint problems . . . 349
20. IKE daemon trace options . . . 352
21. Common NSS server initialization problems . . . 357
22. Common NSS client connection problems . . . 358
23. Common NSS client request failures . . . 360
24. NSS XMLAppliance client API return codes and reason codes . . . 360
25. NSS IPSec client API return codes and reason codes . . . 368
26. NSS server trace options . . . 378
27. SQL problems generating 55x replies (FTP Server) . . . 448
28. Other SQL problems (FTP Server) . . . 449
29. SQL problems generating 55x replies (FTP Client) . . . 481
30. Other SQL problems (FTP Client) . . . 482
31. Debug trace options . . . 502
32. Telnet login problems . . . 517
33. Incorrect output types for Telnet . . . 520
34. Telnet commands from RFC 854 . . . 532
35. Telnet command options from RFC 1060 . . . 533
36. Format of Record 1 of an SMTP ADDRBLOK data set . . . 542
37. Format of Record 2 (for an unresolved from record) of an SMTP ADDRBLOK data set . . . 543
38. Format of Record 2 (for a resolved from record) of an SMTP ADDRBLOK data set . . . 543
39. Format of Record 3 (for an unresolved from record) of an SMTP ADDRBLOK data set . . . 544
40. Debugging switches by category . . . 547
41. qf File code letters . . . 550
42. Format of a SNALINK trace table entry . . . 562
43. Common SNALINK LU6.2 address space startup problems . . . 569
44. Common DLC connection problems . . . 572
45. SNALINK LU6.2 address space network establishment problems . . . 575
46. Outage problem . . . 577
47. Common SNALINK LU6.2 data loss problems . . . 578
48. Common Policy Agent initialization problems . . . 662
49. Policy definition problems . . . 664
50. Common policy client connection problems . . . 668
51. Common policy client retrieval problems . . . 672
52. Common import requestor connection problems . . . 674
53. Common import requestor retrieval problems . . . 678
54. LDAP object retrieval problems . . . 678
55. LDAP object storage problems . . . 680

© Copyright IBM Corp. 2000, 2015

56. Common RSVP initialization problems . . . 688
57. RSVP application problems . . . 689
58. RSVP policy problems . . . 690
59. AT-TLS trace levels . . . 709
60. Common System SSL return codes . . . 712
61. AT-TLS return codes . . . 714
62. SIOCTTLSCTL error codes . . . 718
63. Common defense manager daemon (DMD) problems . . . 722
64. Defense manager daemon (DMD) trace options . . . 728
65. GUI-level object mapping . . . 733
66. IPSec messages logged by TRMD . . . 737
67. Defensive filter messages logged by TRMD . . . 738
68. **ipsec -f** display command options . . . 754
69. **ipsec -F** display command options . . . 755
70. **ipsec -F update** or **delete** command options . . . 755
71. **pasearch** commands . . . 757
72. OMPROUTE command terms . . . 763
73. Neighbor event codes . . . 765
74. Neighbor states . . . 767
75. OMPTRACE options . . . 788
76. Types of RIP messages . . . 795
77. MIB variables registered for use by NCPROUTE . . . 796
78. NCPROUTE connection problems . . . 798
79. NCPROUTE command terms . . . 801
80. NCPROUTE routing failures . . . 802
81. NCPROUTE incorrect output . . . 804
82. Symptoms of session outages . . . 805
83. Component problems . . . 832
84. Connection problems . . . 833
85. Error message and return code problems . . . 837
86. Socket data protocol problems . . . 838
87. IMS transaction build problems . . . 842
88. IMS database problems . . . 842
89. Message EZZ9308E name server resolution, part 1 . . . 875
90. Message EZZ9308E name server resolution, part 2 . . . 876
91. Message EZZ9308E name server resolution, part 3 . . . 879
92. Message EZZ9311E name server resolution, part 1 . . . 880
93. Message EZZ9311E name server resolution, part 2 . . . 881
94. Message EZZ9311E name server resolution, part 3 . . . 883
95. Common initialization and logging problems . . . 916
96. Problems using SMTPNOTE CLIST to create mail messages on the JES spool data set . . . 917
97. Diagnosing and resolving Resolver problems . . . 918
98. JES spool data set problems . . . 919
99. Common messages indicating target server status . . . 923
100. Common mail text problems . . . 927
101. Dead letters problems . . . 929
102. Mail Directory problems . . . 930
103. Common MODIFY command . . . 933
104. Common messages indicating checkpoint status . . . 933
105. Common messages used for monitoring resources . . . 935
106. Commands for various types of dumps . . . 951
107. Relationship between RC field and maximum I-field value . . . 959
108. Vendor ID . . . 993

About this document

This document tells you how to diagnose and report problems occurring in the IBM® z/OS® TCP/IP. Additional information is provided for diagnosing problems with selected applications that are part of z/OS Communications Server.

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.

Use this document to perform the following tasks:

- Diagnose and solve problems in a z/OS Communications Server installation.
- Describe problems to the IBM Software Support Center and document the problems appropriately.

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.

Who should read this document

System programmers can use this document to diagnose problems with TCP/IP or to diagnose problems with z/OS Communications Server components.

How this document is organized

z/OS Communications Server: IP Diagnosis Guide is divided into the following parts:

Part 1, “General diagnosis information,” on page 1 describes how to diagnose a problem suspected to be caused by z/OS Communications Server, select diagnostic tools, and apply diagnostic techniques.

Part 2, “Traces and control blocks,” on page 43 describes selected procedures for TCP/IP Services component trace, packet trace, Socket API trace, and the subcommands (installation, entering, and execution).

Part 3, “Diagnosing z/OS Communications Server components,” on page 313 gives detailed diagnostic information for z/OS Communications Server components.

Appendixes in part 4 provide extra information for this document.

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.

© Copyright IBM Corp. 2000, 2015

This book contains various traces and code examples. In many cases, these examples contain non-release specific information; they are included for illustrative purposes. Actual examples and traces depend on your environment.

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 2 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 website: http://www.software.ibm.com/network/commserver/support/

Most problems can be resolved at this website, where you can submit questions and problem reports electronically, and 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 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 1057.

Conventions and terminology that are 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).

All the exit routines described in this document are *installation-wide exit routines*. The installation-wide exit routines also called installation-wide exits, exit routines, and exits throughout this document.

The TPF logon manager, although included with VTAM®, is an application program; therefore, the logon manager is documented separately from VTAM.

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

Note: In this information, you might see the following Shared Memory Communications over Remote Direct Memory Access (SMC-R) terminology:

- RDMA network interface card (RNIC), which is used to refer to the IBM 10GbE RoCE Express® feature.
- Shared RoCE environment, which means that the 10GbE RoCE Express feature operates on an IBM z13 (z13) or later system, and that the feature can be used concurrently, or shared, by multiple operating system instances. The RoCE Express feature is considered to operate in a shared RoCE environment even if you use it with a single operating system instance.

For definitions of the terms and abbreviations that are 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 that do not have their own syntax described elsewhere.

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 PUt subcommand is a keyword. The required variable parameter is *local_file*, and the optional variable parameter is *foreign_file*. 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.

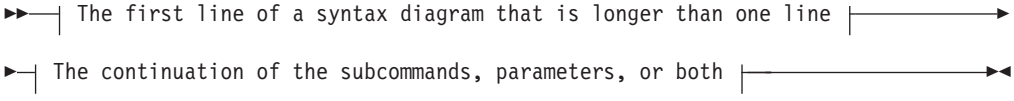

Required operands

Required operands and values appear on the main path line. You must code required operands and values.

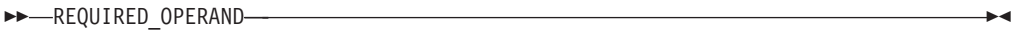

Optional values

Optional operands and values appear below the main path line. You do not have to code optional operands and values.

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.

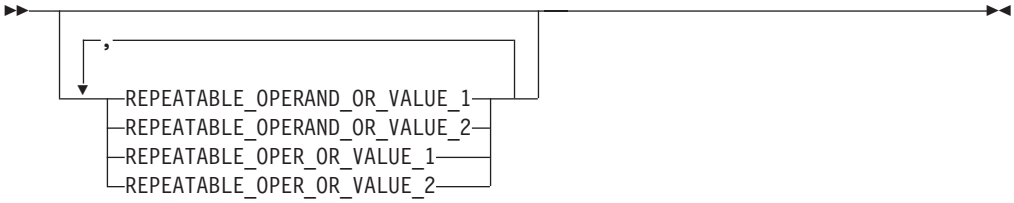

Nonalphanumeric characters

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).

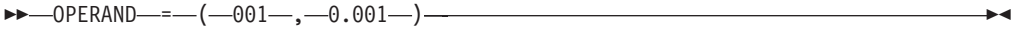

Blank spaces in syntax diagrams

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:

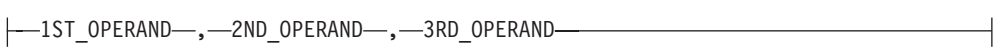

Prerequisite and related information

z/OS Communications Server function is described in the z/OS Communications Server library. Descriptions of those documents are listed in “Bibliography” on page 1041, 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.

Softcopy information

Softcopy publications are available in the following collection.

Titles	Order Number	Description
IBM System z Redbooks Collection	SK3T-7876	The IBM Redbooks® publications selected for this CD series are taken from the IBM Redbooks inventory of over 800 books. All the Redbooks publications that are of interest to the System z® platform professional are identified by their authors and are included in this collection. The System z subject areas range from e-business application development and enablement to hardware, networking, Linux, solutions, security, parallel sysplex, and many others. For more information about the Redbooks publications, see http://www-03.ibm.com/systems/z/os/zos/zfavorites/.

Other documents

This information explains how z/OS references information in other documents.

When possible, this information uses cross-document links that go directly to the topic in reference using shortened versions of the document title. For complete titles and order numbers of the documents for all products that are part of z/OS, see z/OS Information Roadmap (SA23-2299). The Roadmap describes what level of documents are supplied with each release of z/OS Communications Server, and also describes each z/OS publication.

To find the complete z/OS library, visit the z/OS library in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSLTBW/welcome).

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
DNS and BIND, Fifth Edition, O'Reilly Media, 2006	ISBN 13: 978-0596100575
Routing in the Internet, Second Edition, Christian Huitema (Prentice Hall 1999)	ISBN 13: 978-0130226471
sendmail, Fourth Edition, Bryan Costales, Claus Assmann, George Jansen, and Gregory Shapiro, O'Reilly Media, 2007	ISBN 13: 978-0596510299
SNA Formats	GA27-3136
TCP/IP Illustrated, Volume 1: The Protocols, W. Richard Stevens, Addison-Wesley Professional, 1994	ISBN 13: 978-0201633467
TCP/IP Illustrated, Volume 2: The Implementation, Gary R. Wright and W. Richard Stevens, Addison-Wesley Professional, 1995	ISBN 13: 978-0201633542
TCP/IP Illustrated, Volume 3: TCP for Transactions, HTTP, NNTP, and the UNIX Domain Protocols, W. Richard Stevens, Addison-Wesley Professional, 1996	ISBN 13: 978-0201634952
TCP/IP Tutorial and Technical Overview	GG24-3376
Understanding LDAP	SG24-4986
z/OS Cryptographic Services System SSL Programming	*SC14-7495*
z/OS IBM Tivoli Directory Server Administration and Use for z/OS	*SC23-6788*

Title	Number
z/OS JES2 Initialization and Tuning Guide	*SA32-0991*
z/OS Problem Management	*SC23-6844*
z/OS MVS Diagnosis: Reference	*GA32-0904*
z/OS MVS Diagnosis: Tools and Service Aids	*GA32-0905*
z/OS MVS Using the Subsystem Interface	*SA38-0679*
z/OS Program Directory	*GI11-9848*
z/OS UNIX System Services Command Reference	*SA23-2280*
z/OS UNIX System Services Planning	*GA32-0884*
z/OS UNIX System Services Programming: Assembler Callable Services Reference	*SA23-2281*
z/OS UNIX System Services User's Guide	*SA23-2279*
z/OS XL C/C++ Runtime Library Reference	*SC14-7314*
zEnterprise System and System z10 OSA-Express Customer's Guide and Reference	*SA22-7935*

Redbooks publications

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

Title	Number
IBM z/OS V2R1 Communications Server TCP/IP Implementation, Volume 1: Base Functions, Connectivity, and Routing	SG24-8096
IBM z/OS V2R1 Communications Server TCP/IP Implementation, Volume 2: Standard Applications	SG24-8097
IBM z/OS V2R1 Communications Server TCP/IP Implementation, Volume 3: High Availability, Scalability, and Performance	SG24-8098
IBM z/OS V2R1 Communications Server TCP/IP Implementation, Volume 4: Security and Policy-Based Networking	SG24-8099
IBM Communication Controller Migration Guide	SG24-6298
IP Network Design Guide	SG24-2580
Managing OS/390 TCP/IP with SNMP	SG24-5866
Migrating Subarea Networks to an IP Infrastructure Using Enterprise Extender	SG24-5957
SecureWay Communications Server for OS/390 V2R8 TCP/IP: Guide to Enhancements	SG24-5631
SNA and TCP/IP Integration	SG24-5291
TCP/IP in a Sysplex	SG24-5235
TCP/IP Tutorial and Technical Overview	GG24-3376
Threadsafe Considerations for CICS	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/systems/z/os/zos/

z/OS Internet Library

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

www.ibm.com/systems/z/os/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 Communications Server performance information

This site contains links to the most recent Communications Server performance reports.

http://www.ibm.com/support/docview.wss?uid=swg27005524

IBM Systems Center publications

Use this site to view and order Redbooks publications, 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/atsmastr.nsf

Tivoli® NetView® for z/OS

Use this site to view and download product documentation about Tivoli NetView for z/OS

http://www.ibm.com/support/knowledgecenter/SSZJDU/welcome

RFCs

Search for and view Request for Comments documents in this section of the Internet Engineering Task Force website, 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 website

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

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

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

DNS websites

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

USENET news groups
comp.protocols.dns.bind

BIND mailing lists
https://lists.isc.org/mailman/listinfo

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.

The z/OS Basic Skills Information Center

The z/OS Basic Skills Information Center is a web-based information resource intended to help users learn the basic concepts of z/OS, the operating system that runs most of the IBM mainframe computers in use today. The Information Center is designed to introduce a new generation of Information Technology professionals to basic concepts and help them prepare for a career as a z/OS professional, such as a z/OS systems programmer.

Specifically, the z/OS Basic Skills Information Center is intended to achieve the following objectives:

- Provide basic education and information about z/OS without charge
- Shorten the time it takes for people to become productive on the mainframe
- Make it easier for new people to learn z/OS

To access the z/OS Basic Skills Information Center, open your web browser to the following website, which is available to all users (no login required):
http://www-01.ibm.com/support/knowledgecenter/zosbasics/com.ibm.zos.zbasics/homepage.html?cp=zosbasics%2F0

Summary of changes for IP Diagnosis Guide

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

Changes made in z/OS Version 2 Release 2

This document contains information previously presented in z/OS Communications Server: IP Diagnosis Guide, GC27-3652-02, which supported z/OS Version 2 Release 1.

New information

- Enhanced IKED scalability, see "Resource constraint problems" on page 349.
- Activate Resolver trace without restarting applications, see the following topics:
  - "CTRACE TRACERES" on page 909
  - "Enabling collection of CTRACE TRACERES records" on page 910
  - "Disabling collection of CTRACE TRACERES records" on page 910
  - "Formatting the CTRACE TRACERES option" on page 911
- CSSMTP migration enablement, see the following topics:
  - "Duplicate mail messages" on page 932
  - "Verifying mail workload compatibility with SMTPD" on page 932

Changed information

- 64-bit enablement of the TCP/IP stack, see the following topics:
  - "Dumps" on page 13
  - "Sample output of the TCPIPCS TSDB subcommand" on page 264
  - "Limiting CSM storage utilization" on page 943
- Removed support for the GATEWAY statement in the TCP/IP profile, see the following topics:
  - "Netstat ROUTE/-r" on page 39
  - "Sample output of the TCPIPCS PROFILE subcommand" on page 212
  - "Common configuration mistakes" on page 567
  - "Steps for NCPROUTE connection problems" on page 798
  - "Operation" on page 818
  - "Sources of diagnostic information" on page 820
  - "Session hangs" on page 824
  - "Maximum transmission unit (MTU)" on page 956
  - "IPv4 subnetting" on page 966
  - "Static routing" on page 969
- Reordering of cached Resolver results, see the following topics:
  - "Sample output of the RESOLVER command" on page 283
  - "TRACE RESOLVER" on page 884
  - "Interpreting the Trace Resolver output" on page 889

© Copyright IBM Corp. 2000, 2015

- Enhanced IKED scalability, see "Obtaining syslog debug information for the IKE daemon" on page 351.
- Activate Resolver trace without restarting applications, see the following topics:
  - "TRACE RESOLVER" on page 884
  - "Interpreting the Trace Resolver output" on page 889
  - "CTRACE - RESOLVER" on page 907
- CSSMTP migration enablement, see Appendix C, "IBM Health Checker for z/OS," on page 999.

Changes made in z/OS Version 2 Release 1, as updated September 2014

This document contains information previously presented in z/OS Communications Server: IP Diagnosis Guide, GC27-3652-01, which supported z/OS Version 2 Release 1.

Changes made in z/OS Version 2 Release 1, as updated December 2013

This document contains information previously presented in z/OS Communications Server: IP Diagnosis Guide, GC27-3652-00, which supported z/OS Version 2 Release 1.

Changed information

- Network security enhancements for SNMP, see the following topics:
  - "No response from the SNMP agent" on page 627
  - "SNMPv3 AES encryption or decryption failures" on page 629
- TLS security enhancements for policy agent, see "Import requestor connection problems" on page 673.
- AT-TLS enablement for DCAS, see "Analyzing start problems with the DCAS" on page 866.

Summary of changes for z/OS Version 2 Release 1

For specifics on the enhancements for z/OS Version 2, Release 1, see the following publications:

- z/OS Summary of Message and Interface Changes
- z/OS Introduction and Release Guide
- z/OS Planning for Installation
- z/OS Migration

Part 1. General diagnosis information

© Copyright IBM Corp. 2000, 2015

Chapter 1. Overview of diagnosis procedure

To diagnose a problem that is suspected to be caused by z/OS Communications Server, first identify the problem, then determine whether it is a problem with TCP/IP. If the problem is related to TCP/IP, gather information about the problem so that you can report the source of the problem to the IBM Software Support Center. With this information, you can work with IBM Software Support Center representatives to solve the problem.

Steps for diagnosing problems

These steps help you identify the source of a problem and obtain a solution.

About this task

Figure 1 on page 4 summarizes the procedure to follow to diagnose a problem. The steps that follow the figure provide more information about this procedure.

© Copyright IBM Corp. 2000, 2015

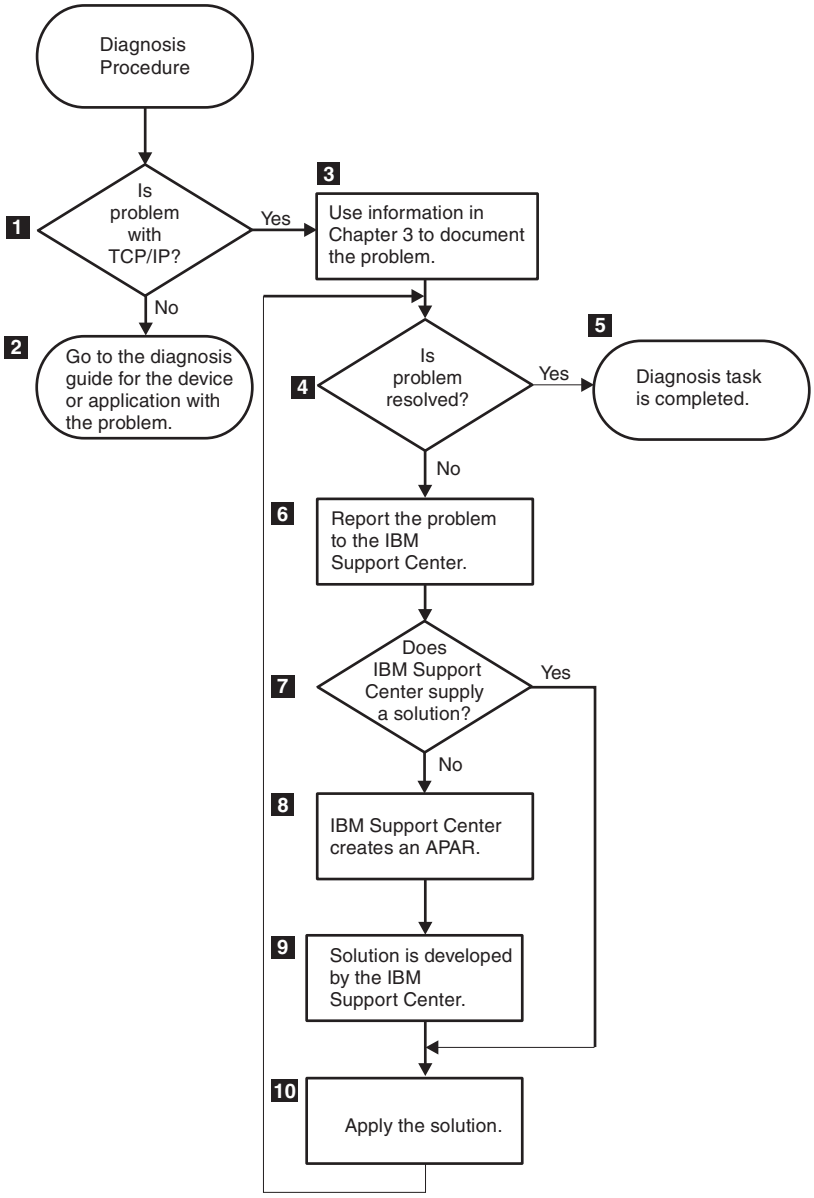

Figure 1. Overview of the diagnosis procedure

Procedure

Perform the following steps to diagnosis a problem:

1. Check sources for diagnostic information.

   Various messages appearing in the console log or in the SYSPRINT or SYSERROR data sets, together with alerts and diagnostic aids, provide information that helps you to find the source of a problem. You should also check syslogd output, and syslog daemon messages, and be prepared to provide this information to the IBM Software Support Center. If the problem is with TCP/IP, go to Step 3; otherwise, go to Step 2.
2. Check appropriate books.

   See the diagnosis guide of the hardware device or software application that has the problem.
3. Gather information.

See Chapter 2, "Selecting tools and service aids," on page 7, for a detailed explanation of diagnostic procedures and how to collect information relevant to the problem.

4. Try to solve the problem.

   If you cannot solve the problem, go to Step 6.

5. The diagnosis task is completed.

   The problem has been solved.

6. Report the problem to the IBM Software Support Center.

   After you have gathered the information that describes the problem, report it to the IBM Software Support Center. If you are an IBMLink user, you can perform your own RETAIN® searches to help identify problems. Otherwise, a representative uses your information to build keywords to search the RETAIN database for a solution to the problem.

   Alternatively, go to http://www.ibm.com/software/network/commserver/support.

   The object of this keyword search using RETAIN is to find a solution by matching the problem with a previously reported problem. When IBM develops a solution for a new problem, it is entered into RETAIN with a description of the problem.

7. Work with IBM Support Center representatives.

   If a keyword search matches a previously reported problem, its solution might also correct this problem. If so, go to Step 10. If a solution to the problem is not found in the RETAIN database, the IBM Software Support Center representatives continue to work with you to solve the problem. Go to Step 8.

8. Create an APAR.

   If the IBM Software Support Center does not find a solution, they create an authorized program analysis report (APAR) in the RETAIN database.

9. A solution is developed by the IBM Software Support Center.

   Using information supplied in the APAR, IBM Software Support Center representatives determine the cause of the problem and develop a solution for it.

10. Apply the solution.

    Apply the corrective procedure supplied by the IBM Software Support Center to correct the problem.

Results

Go to Step 4 to verify that the problem is corrected. You know that you are done when the problem is corrected.

Chapter 2. Selecting tools and service aids

This topic introduces the tools and service aids that z/OS Communications Server provides for diagnosis. As used in this document, the term *tools* includes dumps and traces, while the term *service aids* includes all other facilities provided for diagnosis.

For example:

- SVC dump and system trace are tools.
- LOGREC data set and IPCS are service aids.

The following information is discussed in this topic:

- "How do I know which tool or service aid to select?" lists problem types and matches them with the appropriate tool or service aid. Use this topic to select the tool or service aid you need for a particular problem.
- "Overview of available tools and service aids" on page 13 describes each tool and service aid, including when to use it for diagnosis. Use this topic when you need an overview of tools and service aids, or to find the appropriate time to use a particular tool or service aid.
- "Methods for submitting documentation" on page 17 describes how to send documentation electronically to IBM using FTP or e-mail.
- "Necessary documentation" on page 18 lists the documentation you need to gather before contacting the IBM Software Support Center.

How do I know which tool or service aid to select?

This topic describes the criteria for selecting a tool or service aid.

Your choice depends on one of the following problems or needs:

Problem or need	See
Selecting a dump	"Selecting a dump" on page 8
Selecting a TCP/IP services component trace	"Selecting a trace" on page 8
Selecting a service aid	Table 1

The tables show the problem, the corresponding tool or service aid, and the topic or document that covers it in more detail. Use these tables to find a tool or service aid quickly.

Tip: The traces given in this document are only examples. Traces in your environment can differ from these examples because of different options selected.

Table 1. Selecting a service aid

If the problem is...	Then use this type of service aid
System or hardware problem: need a starting point for diagnosis or diagnosis requires an overview of system and hardware events in chronological order.	LOGREC data set or EREP See z/OS MVS Diagnosis: Tools and Service Aids for detailed information.

© Copyright IBM Corp. 2000, 2015

Table 1. Selecting a service aid (continued)

If the problem is...	Then use this type of service aid
Information about the contents of load modules and program objects or a problem with modules on the system.	AMBLIST See z/OS MVS Diagnosis: Tools and Service Aids for detailed information.
Diagnosis requires a trap to catch problem data while a program is running. The DISPLAY TCPIP,,STOR command can be used to help set a SLIP trap.	Service Level Indication Processing (SLIP) See z/OS MVS System Commands for detailed information.
Diagnosis requires formatted output of problem data, such as a dump or trace.	IPCS See z/OS MVS IPCS User's Guide for detailed information.

You can now perform the steps for the decision you have made.

Selecting a dump

Base your choice of dumps on the criteria given in Table 2.

Table 2. Selecting a dump

If the problem is...	Then use this type of dump
Abnormal end of an authorized program or a problem program.	ABEND dump See "Analyzing abends" on page 21 for detailed information.
TCP/IP server or client address space stops processing or is stopped by the operator because of slowdown or looping condition.	SVC dump The SVC dump is created using the DUMP command. See "Analyzing loops" on page 22 for detailed information.

You can now perform the steps for the decision you have made.

Selecting a trace

Base your choice of traces on the criteria given in Table 3.

Table 3. Selecting a trace

If the problem is...	Then use this type of trace or command	Trace output location
Load balancing using the z/OS Load Balancing Advisor See Chapter 7, "Diagnosing problems with the z/OS Load Balancing Advisor," on page 315 for more information.	Log file	syslogd

Table 3. Selecting a trace (continued)

If the problem is...	Then use this type of trace or command	Trace output location
Network connectivity See Chapter 4, "Diagnosing network connectivity problems," on page 25 for detailed information.	Ping, Netstat ARP/-R For information about Ping, see "Using the Ping command" on page 34. For information about Netstat ARP/-R, see "Netstat ARP/-R" on page 40.	Not applicable
	Packet trace See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for detailed information about packet trace.	CTRACE managed data set
Dynamic VIPA or Sysplex Distributor See Chapter 11, "Diagnosing dynamic VIPA and sysplex problems," on page 383 for detailed information.	Component Trace (SYSTCPIP) XCF option	TCP/IP address space or external writer
TCP/IP socket application See "Socket API traces" on page 74 for detailed information.	Component Trace (SYSTCPIP) SOCKAPI option	TCP/IP address space or external writer
LPR client See "LPR client traces" on page 425 for detailed information.	LPR command with the TRACE option	sysout
LPD server See "LPD server traces" on page 431 for detailed information.	See "LPD server traces" on page 431 for ways to activate traces.	SYSPRINT
z/OS UNIX FTP server See Chapter 14, "Diagnosing File Transfer Protocol (FTP) problems," on page 443 for detailed information.	z/OS UNIX FTP server trace	Server traces appear on the console if syslogd is not started. If it is started, traces appear in the file designated in the syslog.conf file. See z/OS Communications Server: IP Configuration Guide for detailed information about syslogd.
z/OS UNIX Telnet See Chapter 15, "Diagnosing z/OS UNIX Telnet daemon (otelnetd) problems," on page 501, for detailed information.	z/OS UNIX Telnet traces	syslogd
TN3270E Telnet server See Chapter 16, "Diagnosing Telnet problems," on page 515 for detailed information.	Telnet traces	Telnet address space or external writer
SMTP See "SMTP RESOLVER trace" on page 544 for detailed information.	Resolver Trace	Job log output

Table 3. Selecting a trace (continued)

If the problem is...	Then use this type of trace or command	Trace output location
Popper See Chapter 18, "Diagnosing z/OS UNIX sendmail and popper problems," on page 547 for detailed information.	Popper Messages	syslogd
SNALINK LU0 See Chapter 19, "Diagnosing SNALINK LU0 problems," on page 557 for detailed information.	IP Packet Trace	CTRACE managed data set
	Debug Trace	SNALINK LU0 address space
SNALINK LU6.2 See Chapter 20, "Diagnosing SNALINK LU6.2 problems," on page 565 for detailed information.	TRACE DETAIL ALL	SYSPRINT
	IP Packet Trace	CTRACE managed data set
	TCP/IP Internal Trace CTRACE managed data set	CTRACE managed data set
	VTAM Buffer Trace	GTF managed data set, see z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures for detailed information.
Dynamic domain name system (DDNS)	Error messages	syslogd
	Resolver Trace	Job log output
	TCP/IP component trace	CTRACE managed data set
z/OS UNIX REXEC See Chapter 22, "Diagnosing z/OS UNIX REXEC, RSH, REXECD, and RSHD problems," on page 599.	z/OS UNIX REXEC debug trace	syslogd
z/OS UNIX REXECD See Chapter 22, "Diagnosing z/OS UNIX REXEC, RSH, REXECD, and RSHD problems," on page 599.	z/OS UNIX REXECD debug trace	syslogd
z/OS UNIX RSHD See Chapter 22, "Diagnosing z/OS UNIX REXEC, RSH, REXECD, and RSHD problems," on page 599.	z/OS UNIX RSHD debug trace	syslogd
X Window and Motif See Chapter 23, "Diagnosing X Window System and Motif problems," on page 605 for detailed information.	XWTRACE and XWTRACEC (environment variables)	stderr

Table 3. Selecting a trace (continued)

If the problem is...	Then use this type of trace or command	Trace output location
SNMP See Chapter 24, "Diagnosing Simple Network Management Protocol (SNMP) problems," on page 609 for detailed information.	Manager Traces	Console (snmp) or SYSPRINT (NetView SNMP)
	• SNMP Agent Traces • TCP/IP Subagent Traces • OMPROUTE Subagent Traces • Network SLAPM2 Subagent Traces • TN3270E Telnet Subagent Traces • TRAPFWD Traces	syslogd
Policy Agent See Chapter 25, "Diagnosing Policy Agent problems," on page 657 for detailed information.	Log file	See z/OS Communications Server: IP Configuration Guide for detailed information.
RSVP Agent See Chapter 26, "Diagnosing RSVP agent problems," on page 685 for detailed information.	Log file	See z/OS Communications Server: IP Configuration Guide for detailed information.
Traffic Regulator Management Daemon (TRMD) See Chapter 27, "Diagnosing intrusion detection problems," on page 699 for detailed information.	Log file	syslogd
IKE daemon See Chapter 9, "Diagnosing IKE daemon problems," on page 333.	Component trace For detailed information about IKE daemon component trace, see "TCP/IP services component trace for the IKE daemon" on page 351.	CTRACE managed data set
	Log file For detailed information about obtaining IKE daemon debug log information, see "Obtaining syslog debug information for the IKE daemon" on page 351.	syslogd
Network security services (NSS) server See Chapter 10, "Diagnosing network security services (NSS) server problems," on page 357.	Component trace For detailed information about network security services (NSS) server component trace, see "TCP/IP services component trace for the network security services (NSS) server" on page 377.	CTRACE managed data set
	Log file For detailed information about obtaining network security services (NSS) server debug log information, see "Obtaining syslog debug information for the network security service server" on page 367.	syslogd

Table 3. Selecting a trace (continued)

If the problem is...	Then use this type of trace or command	Trace output location
OMPROUTE See Chapter 31, "Diagnosing OMPROUTE problems," on page 759.	Component trace For detailed information about OMPROUTE Component Trace, see "TCP/IP services component trace for OMPROUTE" on page 786.	CTRACE managed data set
	OMPROUTE Trace For detailed information, see "OMPROUTE traces and debug information" on page 775.	stdout
NCPROUTE See Chapter 32, "Diagnosing NCPROUTE problems," on page 793 for detailed information.	NCPROUTE Traces	SYSPRINT
X.25 NPSI See Chapter 33, "Diagnosing X.25 NPSI problems," on page 817 for detailed information.	Server activity log	SYSPRINT
IMS™ See Chapter 34, "Diagnosing IMS problems," on page 827 for detailed information.	IP Packet Trace	CTRACE managed data set
	TCP/IP Internal Trace	CTRACE managed data set
	IMS Trace	For more information, see the IMS knowledge center at http://www-01.ibm.com/support/knowledgecenter/SSEPH2/welcome.
CICS® See Chapter 36, "Diagnosing problems with IP CICS sockets," on page 855 for detailed information.	CICS external trace data set (auxtrace)	For more information, see the CICS knowledge center at http://www-01.ibm.com/software/htp/cics/library/.
	TCP/IP Internal trace	CTRACE managed data set
Express Logon See Chapter 38, "Diagnosing problems with Express Logon," on page 865 for detailed information.	Log file	syslogd
Resolver See Chapter 39, "Diagnosing resolver problems," on page 869 for detailed information.	Trace Resolver	SYSPRINT or stdout
	Resolver Internal trace	CTRACE managed data set

You can now perform the steps for the decision you have made.

Selecting a service aid

Base your choice of service aid on the criteria given in Table 1 on page 7.

Overview of available tools and service aids

This topic provides an overview of the tools and service aids in detail. The topics that follow contain a brief description of each tool or service aid, reasons why you would use it, and a reference to the topic or document that covers the tool or service aid in detail. (Most of the detailed information about tools and service aids is in this document.)

A description of tools and service aids are included in the following sections:

- Dumps, see Table 4
- Traces, see Table 5 on page 14
- Display commands, see "Display commands" on page 16
- System service aids, see Table 6 on page 16

In the tables that follow, the dumps, traces, or service aids are listed by frequency of use.

Tip: The traces given in this document are only examples. Traces in your environment can differ from these examples because of different options selected.

Dumps

Table 4 describes the types of available dumps.

Table 4. Description of dumps

Type of dump	Description
ABEND dumps	Use an ABEND dump when ending an authorized program or a problem program because of an uncorrectable error. These dumps show: • The virtual storage for the program requesting the dump. • System data associated with the program. The system can produce three types of ABEND dumps— SYSABEND, SYSMDUMP, and SYSUDUMP. Each one dumps different areas. Select the dump that gives the areas needed for diagnosing your problem. The IBM-supplied defaults for each dump are: • SYSABEND dumps. The largest of the ABEND dumps, containing a summary dump for the failing program plus many other areas useful for analyzing processing in the failing program. • SYSMDUMP dumps. Contains a summary dump for the failing program, plus some system data for the failing task. In most cases, SYSMDUMP dumps are recommended, because they are the only ABEND dumps that are formatted with IPCS. • SYSUDUMP dumps. The smallest of the ABEND dumps, containing only data and areas about the failing program. See z/OS MVS Diagnosis: Tools and Service Aids for more information about ABEND.

Table 4. Description of dumps (continued)

Type of dump	Description
SVC dumps	SVC dumps can be used in two different ways: • Most commonly, a system component requests an SVC dump when an unexpected system error occurs, but the system can continue processing. • An authorized program or the operator can also request an SVC dump when diagnostic data is needed to solve a problem. SVC dumps contain a summary dump, control blocks, and other system code, but the exact areas dumped depend on whether the dump was requested by a macro, command, or SLIP trap. SVC dumps can be analyzed using IPCS. See z/OS MVS Diagnosis: Tools and Service Aids for detailed information. If a console dump or SLIP is requested: • Capture the OMVS and (if applicable) affected application address spaces as well as TCP/IP. • SDATA specification should contain the RGN, TRT, PSA, SUM, CSA and SQA keywords (at minimum).
Stand-alone dumps	Use a stand-alone dump when: • The system stops processing. • The system enters a wait state with or without a wait state code. • The system enters an instruction loop. • The system is processing slowly. These dumps show central storage and some paged-out virtual storage occupied by the system or stand-alone dump program that failed. Stand-alone dumps can be analyzed using IPCS. See "Analyzing loops" on page 22 for detailed information.

Traces

Table 5 describes the types of available traces.

Table 5. Description of traces

Type of trace	Description
Component trace	Use a component trace when you need trace data to report a client/server component problem to the IBM Software Support Center. Component tracing shows processing between the client and server. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for detailed information.
Data trace	Use a data trace to trace socket data (transforms) into and out of the physical file structure (PFS). See "Data trace (SYSTCPDA) for TCP/IP stacks" on page 146 for detailed information.

Table 5. Description of traces (continued)

Type of trace	Description
GTF trace	Use a Generalized Trace Facility (GTF) trace to show system processing through events occurring in the system over time. The installation controls which events are traced. Use GTF when you are familiar enough with the problem to pinpoint the one or two events required to diagnose your system problem. GTF can be run to an external data set. See z/OS MVS Diagnosis: Tools and Service Aids for more information about GTF.
Master trace	Use the master trace to show the messages to and from the master console. Master trace is useful because it provides a log of the most recently issued messages. These can be more pertinent to your problem than the messages accompanying the dump itself. You can either accept a dump or write this trace to GTF. See z/OS MVS Diagnosis: Tools and Service Aids for detailed information.
OSAENTA trace	Use an OSA-Express network traffic analysis trace to obtain traces of IP packets flowing from and into TCP/IP on a z/OS Communications Server host. The OSAENTA statement lets you copy IP packets as they enter or leave OSA-Express adapter, and then examine the contents of the copied packets. While the packet trace collects data records that flow over the links, the OSAENTA trace collects data records that flow from the network through the OSA adapter. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for detailed information.
Packet trace	Use a packet trace to obtain traces of IP packets flowing from and into TCP/IP on a z/OS Communications Server host. The PKTTRACE statement lets you copy IP packets as they enter or leave TCP/IP, and then examine the contents of the copied packets. While the component trace function collects event data about TCP/IP internal processing, packet trace collects data records that flow over the links. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for detailed information.
System trace	Use system trace to see system processing through events occurring in the system over time. System tracing is activated at initialization and, typically, runs continuously. It records many system events, with minimal details about each. The events traced are predetermined, except for branch tracing. You can either take a dump or write this trace to GTF. See z/OS MVS Diagnosis: Tools and Service Aids for detailed information.
VTAM trace	z/OS Communications Server uses two VTAM components, CSM and MPC. VTAM traces contain entries for many TCP/IP events, especially I/O, and storage requests. See z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT for detailed information.

Table 5. Description of traces (continued)

Type of trace	Description
z/OS UNIX applications	z/OS UNIX applications send debug and trace output to syslogd. For more information about individual components, such as z/OS UNIX FTP or z/OS UNIX SNMP, see those topics in this manual. ITRACE initiated from TCPIP PROFILE processing See the z/OS Communications Server: IP Configuration Guide for more detailed information about syslogd.

Display commands

Display commands can be useful tools and service aids. This topic provides a brief description of the DISPLAY TCPIP,,STOR command. For detailed information about this command, see z/OS Communications Server: IP System Administrator's Commands.

DISPLAY TCPIP,,STOR

Use the DISPLAY TCPIP,,STOR command to display the location and level of a TCP/IP stack module, which verifies that the load module has the appropriate service level.

System service aids

Table 6 lists the service aids supported by z/OS Communications Server.

Table 6. Description of service aids

Type of service aid	Description
AMBLIST	Use AMBLIST when you need information about the contents of load modules and program objects or you have a problem related to the modules on your system. AMBLIST is a program that provides extensive data about modules in the system, such as a listing of the load modules, map of the CSECTs in a load module or program object, list of modifications in a CSECT, map of modules in the LPA, and a map of the contents of the DAT-on nucleus. See z/OS MVS Diagnosis: Tools and Service Aids for more information about AMBLIST.
Common storage tracking	Use common storage tracking to collect data about requests to obtain or free storage in CSA, ECSA, SQA, and ESQA. This is useful to identify jobs or address spaces using an excessive amount of common storage or ending without freeing storage. Use Resource Measurement Facility* (RMF*) or the IPCS VERBEXIT VSMDATA subcommand to display common storage tracking data. • See z/OS RMF User's Guide for more information about RMF™. • See z/OS MVS Initialization and Tuning Guide for detailed information about requesting common storage tracking. • See the VSM topic in z/OS MVS IPCS User's Guide for information about the IPCS VERBEXIT VSMDATA subcommand.

Table 6. Description of service aids (continued)

Type of service aid	Description
Dump suppression	Dump Suppression allows an installation to control dump analysis and elimination (DAE) processing, which suppresses dumps that it considers unnecessary because they duplicate previously taken dumps. DAE suppresses ABEND dumps that would be written to a SYSMDUMP data set (SYSMDUMPs), Transaction dumps (IEATDUMP), and SVC dumps, when the symptom data of a dump duplicates the symptom data of a dump of the same dump type previously taken. DAE uses the ADYSETxx parmlib member to determine the actions DAE is to perform. **Tip:** Consider the SUPPRESSALL statement in ADYSETxx, if dumps are to be considered for suppression. Do this because the Communications Server IP Recovery Routines do not always specify the VRADAE Key in the SDWA(system diagnostic work area) when requesting a dump. See z/OS MVS Initialization and Tuning Guide for more information about requesting dump suppression.
IPCS	Use IPCS to format and analyze dumps, traces, and other data. IPCS produces reports that can help in diagnosing a problem. Some dumps, such as SNAP, SYSABEND, and SYSUDUMP ABEND dumps, are preformatted and are not formatted using IPCS. See z/OS MVS IPCS User's Guide for detailed information.
LOGREC data set	Use the LOGREC data set as a starting point for problem determination. The system records hardware errors, selected software errors, and selected system conditions in the LOGREC data set. LOGREC information gives you an idea of where to look for a problem, supplies symptom data about the failure, and shows the order in which the errors occurred. See z/OS MVS Diagnosis: Tools and Service Aids for detailed information.
SLIP traps	Use serviceability level indication processing (SLIP) to set a trap to catch problem data. SLIP can intercept program event recording (PER) or error events. When an event that matches a trap occurs, SLIP performs the problem determination action that you specify: • Requesting or suppressing a dump • Writing a trace or a LOGREC data set record • Giving control to a recovery routine • Putting the system in a wait state See the SLIP command in z/OS MVS System Commands for detailed information.

Methods for submitting documentation

You can send documentation to IBM using the following methods:

- File Transfer Protocol (FTP)
- email
- TCP/IP active storage or the location and level of a TCP/IP stack module.

Tip: If you use FTP, compress all dumps and traces with the AMATERSE (MVS terse) program, and send the data in BINARY mode.

Requirement: AMATERSE is a prerequisite for PUTDOC.

To obtain PUTDOC and detailed instructions on its use, go to http://www14.software.ibm.com/webapp/set2/sas/f/zaids3/putdoc/putdoc.html.

Using AMATERSE

AMATERSE is an application that prepares diagnostic materials, such as z/OS dumps and traces, for transmission to IBM and vendor sites. When the materials arrive, AMATERSE also provides a means to create similar data sets to support diagnosis of problems.

If you previously used the TRSMAIN utility, you will see that the following changes were made to prepare AMATERSE for formal inclusion in z/OS:

- AMATERSE is used as the preferred application program name rather than TRSMAIN. TRSMAIN is shipped as an alias entry point to AMATERSE.
- The ddnames INFILE and OUTFILE that were required by the TRSMAIN utility are replaced by SYSUT1 and SYSUT2. When the TRSMAIN entry point of AMATERSE is invoked, ddnames INFILE and OUTFILE remain as the defaults.
- AMATERSE is placed into MIGLIB, a library that is part of the link list. No STEPLIB ddname is needed to invoke AMATERSE.
- You can use AMATERSE, the TRSMAIN utility, and VM terse interchangeably in nearly all cases.

Starting AMATERSE

The following sample JCL can be used to start AMATERSE. Lower case text reflects the data that you must alter.

```
//jobname  JOB ...
   //stepname EXEC PGM=AMATERSE,PARM=aaaaa
   //SYSPRINT DD SYSOUT=*,DCB=(RECFM=FBA,LRECL=133,BLKSIZE=12901)
   //SYSUT1   DD DISP=bbb,DSN=your.input.dataset.name
   //SYSUT2   DD DISP=ccc,DCB=ddd,DSN=your.output.dataset.name
   //            SPACE=space_parameters
```

For more information about how to use AMATERSE and any restrictions on its use, see z/OS MVS Diagnosis: Tools and Service Aids.

Using electronic transfer through email attachments

Smaller documents can be sent as attachments to an email message. This can include cut and paste of user output or downloading of the file to a workstation for inclusion. Displayable text can be downloaded by using ASCII transfer; all others should be processed by the AMATERSE utility and transferred in BINARY. Email systems have limits on how much data can be included, so FTP transfers should be used for any significant amounts. (The IBM mail system limit is 10M.)

Necessary documentation

Before you call the IBM Support Center, have the following information available:

Customer number
: The authorization code that allows you to use the IBM Support Center. Your account name, your TCP/IP license number, and other customer identification should also be available.

Problem number
The problem number previously assigned to the problem. If this is your first call about the problem, the support center representative assigns a number to the problem.

Operating system
The operating system that controls the execution of programs (such as z/OS), include the release level.

Language Environment® runtime library
The release level of the link-edit runtime library is also needed if you are compiling user-written applications written in C or C++.

Component ID
A number that is used to search the database for information specific to TCP/IP. If you do not give this number to the support center representative, the amount of time taken to find a solution to your problem increases.

Release number
A number that uniquely identifies each TCP/IP release.

Table 7 lists the specific information that you should provide to the IBM Support Center.

Table 7. TCP/IP component name and release level

Component name and release level	System maintenance program	Field maintenance identifier/CLC
z/OS Communications Server V2R2	SMP/E	The following identifiers are associated with this stack: • HIP6220 (Communications Server IP) • JIP622K (Communications Server Security Level 3) • JIP622X (Communications Server X11R4 XWindows)

The following are component ID numbers for z/OS Communications Server:

Licensed IBM program
z/OS

Component ID number
5650-ZOS

A complex problem might require you to talk to several people when you report your problem to the IBM Support Center. Therefore, you should keep all the information that you have gathered readily available. You might want to keep the items that are constantly required, such as the TCP/IP component ID, in a file for easy access.

Chapter 3. Diagnosing abends, loops, and hangs

This topic contains information about abends, loops, and hangs, located in the following subtopics:

- "Analyzing abends"
- "Analyzing loops" on page 22
- "Steps for analyzing hangs" on page 23

Analyzing abends

An abend is an abnormal end.

Table 8 describes the types of abends that can occur.

Table 8. Types of abends

Type of abend	Description	Where to find help
User Abends	User abends are generated by C run-time routines. They usually start with U409x.	See z/OS Communications Server: IP and SNA Codes.
Platform abends	Abend 3C5 and abend 4C5 are internal abends generated by TCP/IP. Note the reason code stored in register 15 and check the IBM database for known problems.	See z/OS Communications Server: IP and SNA Codes.
System abends	0C4, 0C1, and 878 are system abends.	See z/OS MVS System Codes.
	0D6/0D4/0C4 abends can occur when an application is removed from VMCF/TNF with the F VMCF/TNF, REMOVE command, or if VMCF is not active when an application or command, which requires it is started or issued.	See z/OS MVS System Codes. Can occur when an application is removed from VMCF/TNF with the F VMCF/TNF, REMOVE command. It can also occur when an application or command, which requires it is started or issued. The following TCP/IP applications and commands will abend if VMCF is not active: • SMTP and LPD servers • TSO HOMETEST, LPQ, LPR, LPRM, LPRSET, TELNET, and TESTSITE commands
CEEDUMPs	Language Environment produces certain types of abends detected for z/OS UNIX applications such as z/OS UNIX Telnet. CEEDUMPs are usually written to the current working directory in the hierarchical file structure.	See z/OS Language Environment Debugging Guide.

A dump is usually produced when TCP/IP or a TCP/IP component address space experiences an abend. If an abend occurs and no dump is taken, the dump files or

© Copyright IBM Corp. 2000, 2015

spools might be full or a SYSMDUMP DD statement might not have been specified in the failing procedure. If TCP/IP or a TCP/IP component was not able to complete the dump, gather a console dump of TCP/IP or the failing TCP/IP component, the external trace data set if available, and system log as soon as possible. Otherwise, you must re-create the abend or wait for it to occur again.

For more information about debugging the abends and the system abends (for example, abends 0C4, 0C1, and 878), see z/OS Problem Management.

Analyzing loops

The following are some indicators of a loop:

- Slow response time
- No response at all
- Inordinately high processor utilization by TCP/IP

Steps for collecting documentation

If processing stops or TCP/IP does not respond to commands, TCP/IP might be in a loop and you can collect documentation to help diagnose the problem.

Procedure

Perform the following steps to collect documentation:

1. Get dump output.

   Enabled

   Get an SVC dump of TCP/IP or the looping TCP/IP component by issuing the DUMP command from the MVS system console, or press the Program Restart key. See the z/OS MVS Diagnosis: Tools and Service Aids for more information about the DUMP command.

   Guidelines: Ensure that the following storage areas are dumped because they might be needed to diagnose the TCP loop:

   - TCP/IP and VTAM address spaces.
   - SDATA options RGN, CSA, LSQA, NUC, PSA, and LPA.
   - CSM data spaces. Add DSPNAME=(1.CSM*) to the DUMP command to include both of them in the dump.

   For examples of the DUMP command, see Chapter 5, "TCP/IP services traces and IPCS support," on page 45 and Chapter 42, "Diagnosing storage abends and storage growth," on page 941.

   Disabled

   If the loop is disabled, the MVS system console is not available for input. Try the following steps:

   - Use a PSW RESTART to terminate a looping task. This process creates a LOGREC entry with a completion code of X'071'. Use the LOGREC record and the RTM work area to locate the failing module. Depending on the PSW bit 32, the last 3 bytes (24-bit mode) or 4 bytes (31-bit mode) contain the address being executed at the time of the dump. Scan the dump output to find the address given in the PSW. For more information about using PSW RESTART, see z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures.

- Take a stand-alone dump. See z/OS MVS Diagnosis: Tools and Service Aids for information about stand-alone dumps.

2. Get the MVS system console log (SYSLOG), the job log from the started procedure, and the LOGREC output.

   The MVS system console log might contain information, such as error messages, that can help you diagnose the problem. Also, print the LOGREC file.

   Use the LOGDATA option to print the in-core LOGREC buffers. See z/OS MVS Diagnosis: Tools and Service Aids or z/OS MVS IPCS Commands for more information about the LOGDATA option.

   Tip: The SYSERROR data set might contain additional information to help you diagnose the problem.

3. Determine whether there are any messages associated with the loop, such as a particular message always preceding the problem, or the same message being issued repeatedly. If so, add the message IDs to your problem documentation.
4. Examine the trace entries using IPCS.

   By examining all of the trace entries in the system trace table, you might be able to determine whether there is a loop. The most obvious loops would be a module or modules getting continual control of the TCP/IP system.

   Use the PSW to determine the names of the modules in the loop. See z/OS MVS IPCS User's Guide for information about using IPCS.

   In the output shown in Figure 2, the CLKC entries indicate an enabled loop. The PSW addresses on the CLKCs identify the looping program. Use the WHERE subcommand to locate the responsible program.

```
02-0029 008E7220   CLKC      078D2600 83A8192C  00001004  00000000

02-0029 008E7220   CLKC      078D2600 83A81934  00001004  00000000

02-0029 008E7220   CLKC      078D2600 83A81930  00001004  00000000

02-0029 008E7220   CLKC      078D2600 83A8192A  00001004  00000000

02-0029 008E7220   CLKC      078D2600 83A81930  00001004  00000000

02-0029 008E7220   CLKC      078D2600 83A81938  00001004  00000000
```

Figure 2. Example of output from the IPCS SYSTRACE command

Steps for analyzing hangs

Determine what is hung and what documentation to collect.

Procedure

Perform the following steps to analyze and collect documentation:

1. Determine whether all TCP/IP processing stopped or only processing with respect to a single device, or something in between. Also determine what, if any, recovery action was taken by the operator or user at the time the hang was encountered. Some information about the activity that immediately preceded the hang might be available on the system log or in application program transaction logs.
2. Determine whether TCP/IP responds to commands, such as Ping or Netstat HOME/-h. If TCP/IP does not respond to these commands, take an SVC dump

of TCP/IP address space and contact the IBM Software Support Center. If TCP/IP does respond to the commands, it is not hung.

3. Determine whether a particular application (such as z/OS UNIX FTP or a user-written application) is hung.

   Take a dump of the OMVS address space, the TCP/IP address space, and the application address space.

Chapter 4. Diagnosing network connectivity problems

This topic describes the diagnosis process for network connectivity problems and contains the following subtopics:

- "Communicating through the correct stack" on page 26
- "Steps for diagnosing problems connecting to a server" on page 26
- "Steps for verifying server operation" on page 27
- "Steps for verifying IP routing to a destination when not using policy-based routing" on page 29
- "Steps for diagnosing problems with IP routing to a destination when using policy-based routing" on page 31
- "Steps for verifying network access" on page 33
- "Tools for diagnosing network connectivity problems" on page 34
- "Documentation for the IBM Support Center" on page 41

Overview

Interconnectivity between network hosts encompasses the physical layer or hardware layer, the protocols such as TCP and IP, the IP security services, and the applications that use the services of TCP and IP. To understand interconnectivity, you should first understand internetworking. For detailed information about internetworking, see Appendix A, "Overview of internetworking," on page 955.

Isolating network problems is an essential step in successful implementation of a network application. This topic introduces commands and techniques you can use to diagnose network connectivity problems.

The following diagnostic commands are available for either the z/OS UNIX environment or the TSO environment:

- Ping
- Netstat
- Traceroute

Netstat reports are also available from the console environment by invoking the DISPLAY TCPIP,,NETSTAT command. For complete descriptions of these commands and examples of their output, see z/OS Communications Server: IP System Administrator's Commands.

When referring to these commands and their options throughout this section, both the TSO and z/OS UNIX shell command options are listed, separated by a slash. For example, the recommendation to use Netstat to view the stack's HOME list of IP addresses appears as "use Netstat HOME/-h."

MVS-style data sets are written in capital letters (for example, *hlq*.TCPIP.DATA). Files names in the z/OS UNIX file system are written in lowercase (for example, /etc/hosts).

Table 9 on page 26 lists the name of the commands in each environment.

© Copyright IBM Corp. 2000, 2015

Table 9. Diagnostic commands

UNIX command	TSO command	Refer to:
ping/oping	PING	"Using the Ping command" on page 34
netstat/onetstat	NETSTAT	"Using the Netstat command" on page 39
traceroute/otracert	TRACERTE	"Using the Traceroute command" on page 41

Note: Do not use the resolver and domain name server functions, which translate symbolic names to IP addresses, when diagnosing network problems. Use the host IP address instead.

Communicating through the correct stack

If you are running multiple stacks, the first question to ask is whether the application is communicating through the correct stack. To identify the stack an application is using, you can look at the keyword TCPIPjobname in the TCPIP.DATA file. An application can also select a stack using the SETIBMOPT socket API.

You can use the Netstat parameter TCP/-p to specify the TCP/IP stack name for which you want Netstat report output. This lets you determine the characteristics of a particular stack.

Using the information provided by Netstat, you can change, if necessary, the *hlq*.PROFILE.TCPIP data set or the application configuration file. Alternatively, the application might need to communicate through another stack.

It is also helpful to understand the search order for configuration information used by z/OS Communications Server. See z/OS Communications Server: IP Configuration Reference, "Understanding search orders of configuration information", for more information.

For more information about running multiple stacks, see z/OS Communications Server: IP Configuration Guide.

Steps for diagnosing problems connecting to a server

Determine the source of the problem in connecting to a server.

Procedure

Perform the following steps:

1. Verify that TCP/IP is running correctly on your host. Use Ping loopback, then Ping one of your home addresses. For information about the Ping command, see z/OS Communications Server: IP System Administrator's Commands.
2. Verify that the server application is operational. See "Steps for verifying server operation" on page 27 for more information.
3. Verify IP routing to the server or the client. If you are not using policy-based routing, see "Steps for verifying IP routing to a destination when not using policy-based routing" on page 29 for more information. Otherwise, see "Steps

for diagnosing problems with IP routing to a destination when using policy-based routing" on page 31 for more information.

4. Use the DISPLAY TCPIP,,NETSTAT,ACCESS,NETWORK command to determine whether network access has been configured on the TCP/IP stack. See z/OS Communications Server: IP System Administrator's Commands for more information about this command. If network access control is enabled, then the server might not be permitted to send or receive data on a socket. See "Steps for verifying network access" on page 33 to determine whether network access controls are impacting the server application.
5. Verify IP security protection for the server. If IP security is enabled, then IP traffic to or from the server might not be permitted to flow. See "Steps for diagnosing IP security problems" on page 734 to determine whether IP security controls are impacting the server application.

Steps for verifying server operation

Determine whether the server is active and is using the correct port.

Before you begin

Identify the job name and port of the server to be verified.

About this task

Figure 3 on page 28 shows the decisions involved for verifying server operation.

Verify Server Operation

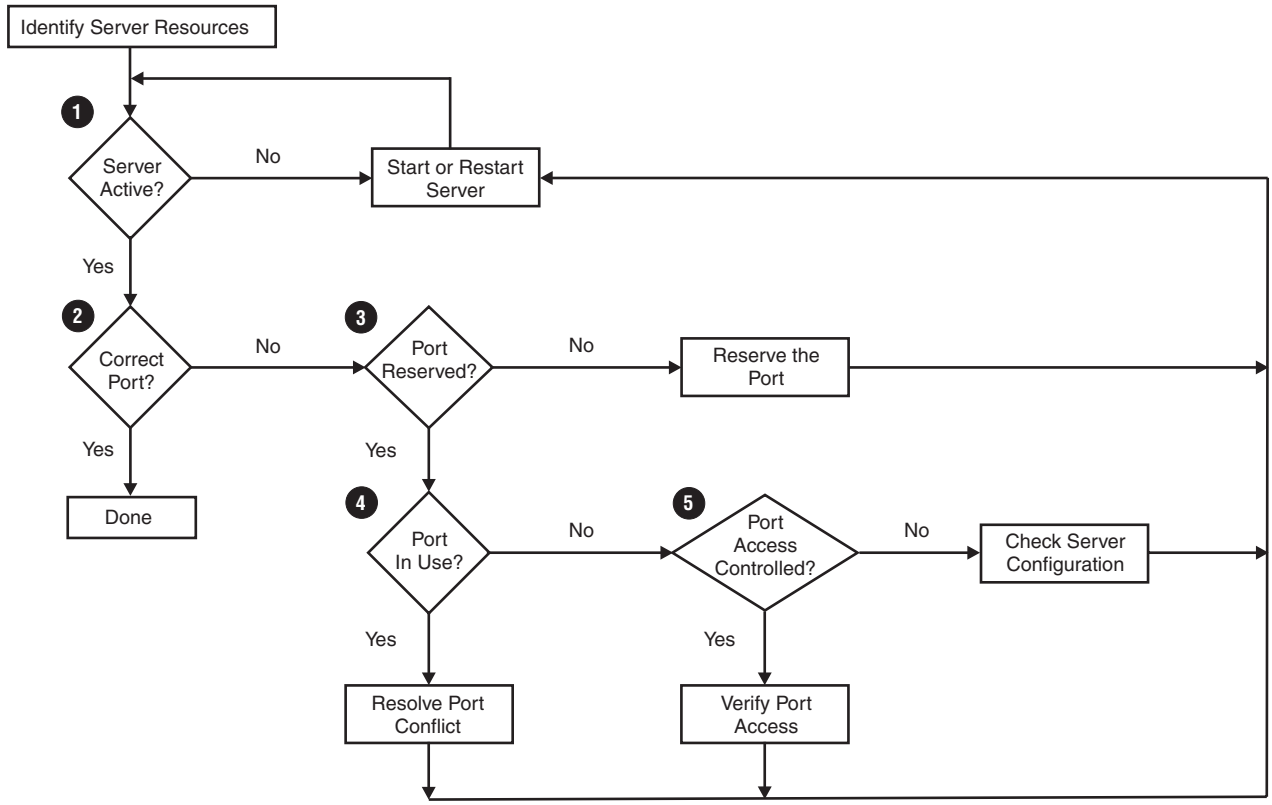

Figure 3. Overview of verifying server operation

Procedure

Perform the following steps to verify server operation:

1. Ensure that the server is started. If not, start the server.
2. Use the Netstat SOCKETS/-s command to determine which port the server is listening on, filtered on the application's job name (-E option for z/OS UNIX, CLIENT keyword for TSO and Operator commands). For example

   ```
   NETSTAT SOCKETS (CLIENT SMTP
   ```

   If the server is not listening on the correct port, configure it correctly. For basic information about the Netstat SOCKETS/-s command, see z/OS Communications Server: IP System Administrator's Commands for details. For details on server configuration, see z/OS Communications Server: IP Configuration Reference.
3. Ensure that there is a PORT statement in the TCP/IP profile data set, to reserve the port for the server. If the server is started but not using the correct port, then a PORT statement might be missing. See z/OS Communications Server: IP Configuration Reference for more information about the PORT statement.
4. Use the Netstat SOCKETS/-s command to determine whether a different server is using the port filtered on the port number (-p option for z/OS UNIX, PORT keyword for TSO and Operator commands). Unless the SHAREPORT keyword is specified on the PORT statement, only one server can be listening on a TCP port. See z/OS Communications Server: IP Configuration Reference for more information about the PORT statement.

5. Check the PORT statement for the server to determine whether the SAF keyword has been specified. If so, then port access control is in effect for the port. See z/OS Communications Server: IP Configuration Guide for more information about port access control. Ensure that the user ID associated with the server is permitted to the security resource name represented by the SAF keyword value. See the description of the PORT statement SAF keyword in the z/OS Communications Server: IP Configuration Reference for information about the security resource name. If the SAF keyword was not specified on the PORT statement, and the server belongs to the z/OS Communications Server product, see z/OS Communications Server: IP Configuration Reference for configuration information that is specific to the server.

Steps for verifying IP routing to a destination when not using policy-based routing

Verify a route to a particular destination IP address when you are not using policy-based routing (PBR).

About this task

Figure 4 on page 30 shows the decisions involved for verifying IP routing to a destination.

Verify IP Routing to a Destination

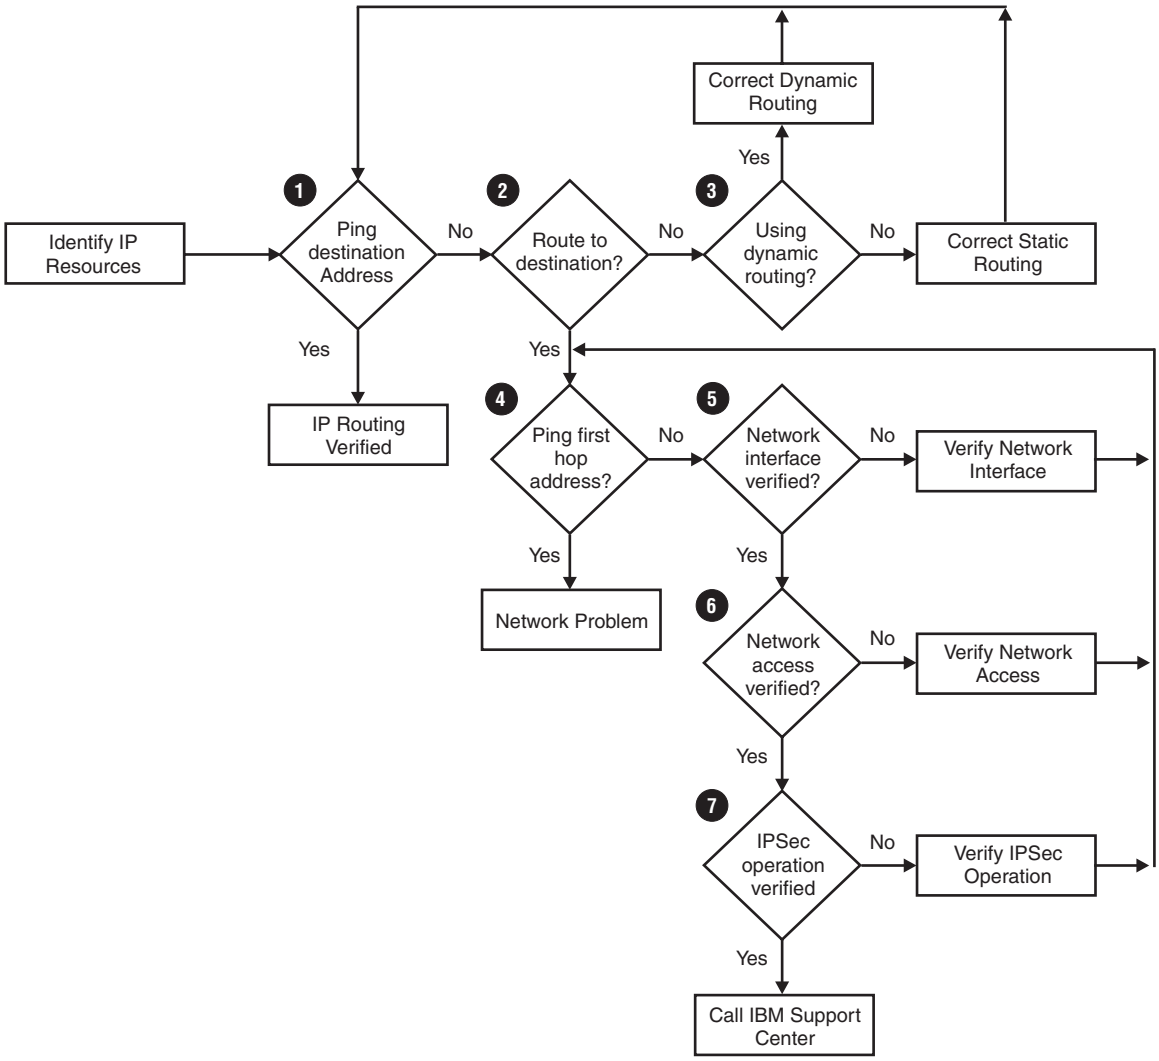

Figure 4. Overview of verifying IP routing to a destination

Procedure

Perform the following steps to verify IP routing to a destination:

1. Use the Ping command to determine whether there is connectivity to the identified IP address. For information about the Ping command, see z/OS Communications Server: IP System Administrator's Commands.
2. If the Ping command fails immediately, there might not be a route to the destination. Use the Netstat ROUTE/-r command to display routes to the network. Verify whether TCP/IP has a route to the destination. For information about the Netstat ROUTE/-r command, see z/OS Communications Server: IP System Administrator's Commands.

   If there is no route, proceed to step 3. If a route exists, proceed to step 4 on page 31.
3. If there is no route to the destination, problem resolution depends on whether static or dynamic routing is being used. See z/OS Communications Server: IP Configuration Guide for more information about static and dynamic routing.

4. If a route exists, verify that the route is correct for the destination. If multipath routing is in effect for the destination, use the Ping command to determine whether there is connectivity to the IP address over any route. Invoke the Netstat CONFIG/-f command and check the value in the output report field, MultiPath, to determine whether multipath routing is in effect and what multipath routing is active.

   Determine whether there is a gateway identified for the route to the destination. If there is no gateway, then the destination address is presumed to be directly connected. In this case, proceed to step 5.

   If a gateway is identified for the route, use the Ping command to confirm connectivity to the gateway. Do one of the following:

   - If the gateway responds to a Ping, then there is a network problem at the gateway or beyond. Use the Traceroute command with the final destination address to determine at which hop in the route the failure is occurring. For information about using the Traceroute command, see z/OS Communications Server: IP System Administrator's Commands.
   - If the gateway does not respond to a Ping, proceed to step 5.
5. Determine which network interface is associated with the route to the destination. If the network interface operation has not been verified for this interface, verify it now. See "Steps for verifying network interface operation" on page 33 for more information.
6. Use the DISPLAY TCPIP,,NETSTAT,ACCESS,NETWORK command to see whether network access control is enabled. If it is enabled, see "Steps for verifying network access" on page 33 for more information.
7. Use the Netstat CONFIG/-f command to determine whether IP security is enabled. If the output report field, IpSecurity, contains the value `Yes`, then IP security is enabled. See "Steps for verifying IP security and defensive filter operation" on page 741 for information about how to verify that IP security is correctly configured. If the problem still exists, see "Documentation for the IBM Support Center" on page 41 to determine what problem documentation you need, and then call the IBM Support Center.

Steps for diagnosing problems with IP routing to a destination when using policy-based routing

Verify a route to a particular destination IP address when you are using policy-based routing (PBR).

Procedure

Perform the following steps to diagnose problems with IP routing to a destination when you are using policy-based routing:

1. While the application is active and attempting to connect to the destination, use the Netstat ALL/-A report to determine the policy rule that is assigned to the connection and the route table being used to perform a route lookup. For information about the Netstat command, see z/OS Communications Server: IP System Administrator's Commands.
   - If no policy rule is listed and the connection is not expected to use policy-based routing, see "Steps for verifying IP routing to a destination when not using policy-based routing" on page 29.
   - Continue to the following step if one of the following is true:
     – A policy rule is not listed and the connection is expected to use policy-based routing

- A policy rule is listed and the connection is not expected to use policy-based routing
- A policy rule is listed, but it is not the expected policy rule

• Otherwise, continue with step3.

2. For information about how to map a connection to the correct policy rule, see the 'Policy-based routing' section in z/OS Communications Server: IP Configuration Guide.
3. Use pasearch to find the policy rule and the corresponding action. For information about the **pasearch** command, see "Displaying policy-based networking information" section of z/OS Communications Server: IP System Administrator's Commands. The policy action will list all the possible route tables that can be used for the connection. Perform steps 4 through 6 on each of the route tables listed in the action.
4. Use the Netstat ROUTE/-r PR command to display routes in the route table. Verify whether TCP/IP has a route to the destination/network in the route table. For information about the Netstat ROUTE/-r command, see z/OS Communications Server: IP System Administrator's Commands.
   • If there is no route to the destination/network and no route is expected to be found in the route table, repeat step 4 using the next route table in the policy action.
   • If there is no route to the destination/network and a route was expected in the route table, see z/OS Communications Server: IP Configuration Guide for information about setting up static and dynamic routing for policy-based routing tables.
   • If a route was found, verify that the route is marked active (has the U flag). If the route is not active, see z/OS Communications Server: IP Configuration Guide for information about route states.
   • If an active route is found, verify that the route table name matches the route table name displayed on the Netstat ALL/-A report for the connection. If it does not, continue to step 9. Otherwise, continue to step 5.
5. Determine whether there is a gateway identified for the route to the destination. If there is no gateway, then the destination address is presumed to be directly connected. In this case, proceed to step 6. If a gateway is identified for the route, use the Ping command to confirm connectivity to the gateway.
   • If the gateway responds to a Ping, then there is a network problem at the gateway or beyond.
   • If the gateway does not respond to a Ping, proceed to step 6.
6. Determine which network interface is associated with the route to the destination. If the network interface operation has not been verified for this interface, verify it now. See “Steps for verifying network interface operation” on page 33 for more information.
7. Use the DISPLAY TCPIP,,NETSTAT,ACCESS,NETWORK command to determine whether network access control is enabled. If it is enabled, see “Steps for verifying network access” on page 33 for more information.
8. Use the Netstat CONFIG/-f command to determine whether IP security is enabled. If the output report field IpSecurity contains the value Yes, then IP security is enabled. If it is enabled, see “Steps for verifying IP security and defensive filter operation” on page 741 for information about how to verify that IP security is correctly configured.
9. See “Documentation for the IBM Support Center” on page 41 to determine what problem documentation you need, and then call the IBM Support Center.

Steps for verifying network interface operation

Determine whether a particular network interface is ready.

About this task

Figure 5 shows the decisions involved for verifying network interface operation.

Verify Network Interface Operation

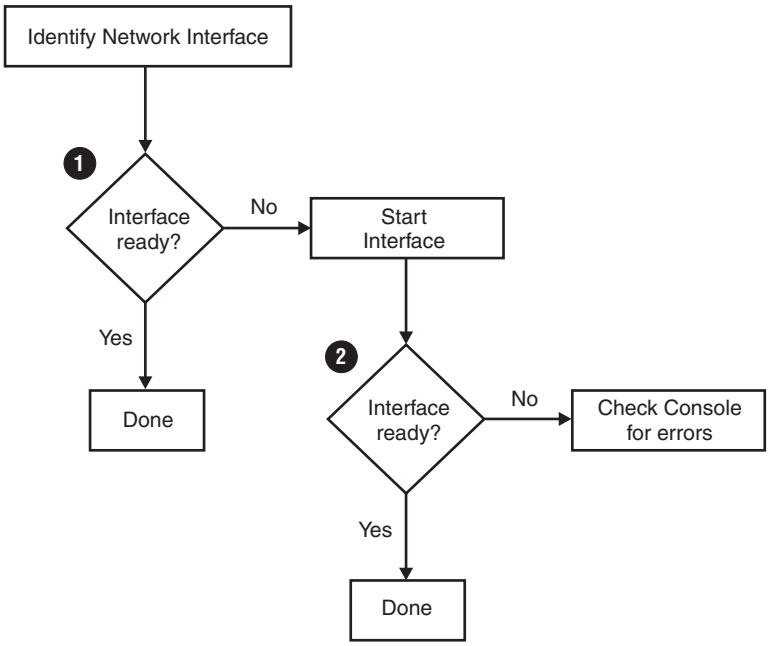

Figure 5. Overview of verifying network interface operation

Procedure

Perform the following steps to verify network interface operation:

1. Use the Netstat DEVLINKS/-d command to check the interface status. If the interface status is Ready, check the physical connectivity from the interface to the network and check for configuration errors in the network. For example, if you are using VLAN, verify that you have configured the proper VLAN IDs throughout the network. If the interface status is not Ready, try to start the interface by using the VARY TCPIP,,START command, and proceed to 2.
2. Use the Netstat DEVLINKS/-d command again to determine whether the interface is ready after being started. If the interface is not ready, check the system console for error messages issued from TCP/IP, VTAM, or IOS and respond as suggested in the documentation for the messages that appear.

Steps for verifying network access

Determine whether access to a resource is controlled and whether the server or client application can access the resource.

Before you begin

Identify the IP address, subnet, or prefix for which network access is to be verified.

About this task

Figure 6 shows the decisions involved for verifying network access.

Verify Network Access

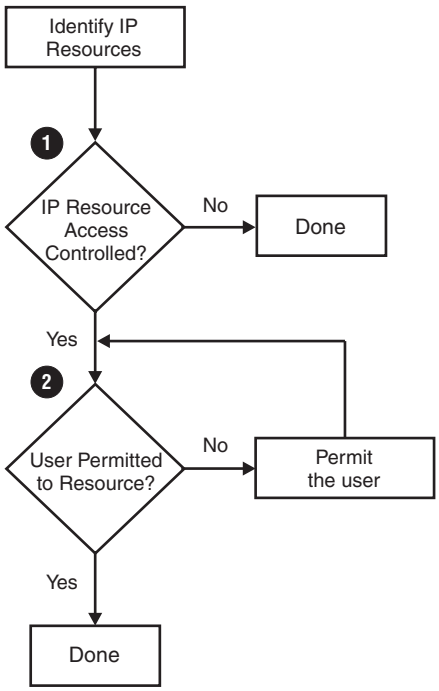

Figure 6. Overview of verifying network access

Procedure

Perform the following steps to verify network access:

1. Invoke the DISPLAY TCPIP,,NETSTAT,ACCESS,NETWORK,ipaddress command, specifying the IP address for which access is to be verified. If the command output indicates that network access control is in effect for the IP address, proceed to 2.
2. Verify that the server or client application is permitted access to the IP resource. See Chapter 12, “Diagnosing access control problems,” on page 417 for more information about verifying this access.

Tools for diagnosing network connectivity problems

This topic describes tools used to diagnose network connection problems.

Using the Ping command

The packet Internet groper (Ping) command sends an Internet Control Message Protocol (ICMP/ICMPv6) Echo Request to a host, gateway, or router with the expectation of receiving a reply. You can invoke the Ping function by using the TSO PING command or the z/OS UNIX shell **ping** or **oping** command.

For a complete description of the Ping command and examples of Ping output, see the z/OS Communications Server: IP System Administrator's Commands.

The Ping command does not use the ICMP/ICMPv6 header sequence number field (icmp_seq or icmp6_seq) to correlate requests with ICMP/ICMPv6 Echo Replies. Instead, it uses the ICMP/ICMPv6 header identifier field (icmp_id or icmp6_id) plus an 8-byte TOD time stamp field to correlate requests with replies. The TOD time stamp is the first 8-bytes of data after the ICMP/ICMPv6 header. When you specify the Verbose/-v parameter, the ICMP/ICMPv6 header sequence numbers sent in the ICMP/ICMPv6 echo requests are displayed in the verbose report of detailed ICMP/ICMPv6 echo replies. Use these sequence numbers to detect the out-of-order and lost packets, based on missing sequence numbers.

When the PMTU/-P parameter with a value of yes or ignore is specified on the command, Ping will ensure that the outbound echo request packets are not fragmented. As a result, ICMP/ICMPv6 error messages may be received by the Ping command if the echo request packet is too large to be sent out by the stack or, forwarded at some point in the network. In this case, the Ping command uses both the ICMP/ICMPv6 header identifier and sequence number fields to correlate requests with the error messages. For IPv6 Ping requests, the Ping command will also use the 8-byte TOD time stamp returned in the ICMPv6 Packet Too Big error message.

Ping can be used in the following ways:

Pinging loopback is used to verify the installation of TCP/IP in the z/OS Communications Server environment.

The Ping loopback is an internal software test. The command examples below use the IPv4 standard loopback address, 127.0.0.1, or the IPv6 standard loopback address, ::1. An IP packet is not sent to a physical device.

```
ping 127.0.0.1
```

For IPv6

```
ping ::1
```

Ping a home address to verify the information from the Netstat HOME/-h command.

This is an internal software test. An IP packet is not sent to a physical device.

```
ping 9.67.113.58
```

Ping a host on a directly attached network to verify the following:

- If equal-cost multipath routes exist in the IP routing table for outbound IP traffic to reach a remote host, use the Ping INTF/-i option to select a routing interface with the attached equal-cost multipath route. Alternatively, for routing interfaces associated with an IPv6 link-local address, the name of the routing interface can be appended as scope information to the IPv6 link-local address of the remote host. When running multiple TCP/IP stacks on the same MVS image, specify the TCP/-p parameter, along with the scope, to indicate the stack to which the routing interface is configured. Whenever applicable, use either of these options to test connectivity. For more information about using scope, see the section on support for scope in z/OS Communications Server: IPv6 Network and Application Design Guide.
- The directly attached network is defined correctly.
- The device is properly connected to the network.
- The device is able to send and receive packets on the network.
- The remote host is able to receive and send packets.

```
ping 9.67.43.101 (intf eth1
ping fe80::9:67:43:104%ethipv61 -p tcpip1
```

Ping a host on a remote network to verify the following:

- If equal-cost multipath routes exist in the IP routing table for outbound IP traffic to reach the remote host, use the Ping INTF/-i option to select a routing interface with the attached equal-cost multipath route. Whenever applicable, use this option to test connectivity.
- The route to the remote network is defined correctly.
- The router is able to forward packets to the remote network.
- The remote host is able to send and receive packets on the network.
- The remote host has a route back to the local host.

```
ping -i eth1 mvs1
```

Note: Ping commands to a remote host might fail if there is a firewall between the two systems, even if the host is reachable using other commands.

Display details of echo replies and obtain summary statistics

You can use the Ping command with the Verbose/-v parameter to obtain detailed echo replies and summary statistics regarding the round-trip times based on the response times in the received echo replies. The detailed echo replies can be used to identify lost echo reply packets based on their sequence numbers and to identify how many hops the echo requests have traveled based on their values of time-to-live (TTL) or maximum number of hops (hop limits). The Verbose/-v parameter provides the following output information:

- Number of bytes for the ICMP data portion
- Echo reply details (for each echo reply received):
  - from: echo reply sender's IP address
  - seq: ICMP/ICMPv6 sequence number
  - ttl: time-to-live for number of hops (if IPv4)
  - hoplim: maximum hop limit (if IPv6)
  - time: response time
- Ping statistics summary:
  - Packets history:
    - Sent: total number of echo requests sent
    - Received: total number of echo replies received
    - Lost: total number of packets lost plus percentage of packet loss
  - Approximate round-trip times in milliseconds:
    - Minimum: minimum value among all RTTs
    - Maximum: maximum value among all RTTs
    - Average: average RTT among all RTTs
    - StdDev for standard deviation among all RTTs

```
ping mvs1 (verbose ping -v mvs1
```

For examples of Ping verbose reports, see z/OS Communications Server: IP System Administrator's Commands.

Determine the path MTU size to a host:

Use the Ping PMTU/-P parameter with the values yes or ignore, to prevent fragmentation of the outbound echo request packets and specify what type

of path MTU discovery support you want. If the outbound packet needs to be fragmented, Ping will display the host name and IP address of the host where the fragmentation is needed.

yes Specifies that the outbound echo request packets will not be fragmented at the local host or in the network, and that you want to use the MTU value, determined by path MTU discovery for the destination.

- If path MTU discovery is enabled and has already determined an MTU value for the destination, and the length of the Ping echo request packet is larger than this MTU size, then the local TCP/IP stack will not send out the packet. In this case, Ping displays one of the local stack's IP addresses as the address of the host where fragmentation is needed, and the next-hop MTU value displayed by Ping is the current path MTU value to the destination. For Ping commands to IPv4 destinations, the Ping command processing will not cause path MTU discovery support to be triggered for the destination. For IPv4, only TCP processing causes path MTU discovery support to be triggered
- If path MTU discovery is not enabled, or has not already determined a path MTU value for the destination, and the Ping echo request packet exceeds the configured route MTU selected for this packet, then the local TCP/IP stack will not send out the packet. In this case, Ping will display one of the local stack's IP addresses as the address of the host where fragmentation is needed, and the next-hop MTU value displayed by Ping is that of the route selected for the Ping packet.
- If the Ping request fails because the echo request packet needs to be fragmented at some point in the network, Ping will display the IP address where fragmentation needs to occur and will display the next-hop MTU value, if it was provided.

ignore Specifies that the outbound echo request packets will not be fragmented at the local host or in the network, and that any MTU values determined by path MTU discovery for the destination, will be ignored.

- If path MTU discovery has determined an MTU value for the destination, and the length of the Ping echo request packet is larger than this MTU size, specifying a value of **ignore** causes the TCP/IP stack to ignore the path MTU value and attempt to send out the packet. As long as the echo request packet length does not exceed the configured route MTU selected for this packet, you can use the ignore value to determine where in the network the original MTU problem occurred. In this case, Ping displays the IP address where fragmentation needs to occur and will display the path MTU value, if it was provided.
- If the Ping echo request packet exceeds the configured route MTU selected for this packet, then the local TCP/IP stack will not send out the packet. In this case, Ping displays one of the local stack's IP addresses as the address of the host where fragmentation is needed, and the next-hop MTU value displayed by Ping is that of the route selected for the Ping packet.

MULTIPATH PERPACKET considerations:

When MULTIPATH PERPACKET is in effect, and equal-cost routes are configured to the Ping destination host, the smallest MTU

value of all the equal-cost routes is used as the largest packet size which can be sent, even if some of the equal-cost routes could support a larger packet size.

- Specify the NONAME/-n parameter to request that Ping only display the IP address of the host, and not attempt to resolve the IP address to a host name. This saves a name server address-to-name lookup. If this host also returned the next-hop MTU size, the size is also displayed.
- Vary the length of the outbound packet to determine where the packet needs to be fragmented. The Length/-l parameter on the Ping command, specifies the number of data bytes for the echo request.
  - For IPv4 destinations, the total length of the outbound echo request packet includes the length of an IPv4 IP header (20 bytes), the length of an ICMP header (8 bytes), and the data length specified by the Length/-l parameter. Depending on your TCP/IP stack configuration, the TCP/IP stack might add additional IP header options to the IP header created by Ping, before the echo request packet is sent, thereby increasing the size of the packet.
  - For IPv6 destinations, the total length of the outbound echo request packet includes the length of an IPv6 IP header (40 bytes), the length of an ICMPv6 header (8 bytes), and the data length specified by the Length/-l parameter. Depending on your TCP/IP stack configuration, the TCP/IP stack might add additional IPv6 extension headers to the packet created by Ping, before the echo request packet is sent, thereby increasing the size of the packet.

Correcting timeout problems

A Ping timeout message can occur for many reasons, and various techniques can be used to identify whether the problem is the local z/OS server or a remote host or router.

Base your actions on the possible reasons for a timeout, as shown in Table 10.

Table 10. Diagnosis of a timeout

If the problem is...	Then use these diagnostic techniques
The device is not transmitting packets to the local network.	Use Netstat DEVLINKS/-d to collect information to help you diagnose the problem. (See DEVLINKS/-d report option in z/OS Communications Server: IP System Administrator's Commands.)
The remote host is not receiving or transmitting packets on the network.	Use Netstat ARP/-R to display the IPv4 entry for the remote host. (See the ARP/-R report option in z/OS Communications Server: IP System Administrator's Commands.) Use Netstat ND/-n to display the IPv6 entry for the remote host. (See the ND/-n report option in z/OS Communications Server: IP System Administrator's Commands.)

Table 10. Diagnosis of a timeout (continued)

If the problem is...	Then use these diagnostic techniques
The remote host does not have a route back to the local z/OS server.	Use Netstat ROUTE/-r on the remote host to make sure that it has a route back. (See ROUTE/-r report option in z/OS Communications Server: IP System Administrator's Commands.)
An intermediate router or gateway is not correctly forwarding IP packets.	Use a packet trace. (See Chapter 5, "TCP/IP services traces and IPCS support," on page 45.)
The IP reassembly timeout value might be set too low.	See the TCP/IP Profile statements, IPCONFIG and IPCONFIG6, in z/OS Communications Server: IP Configuration Reference.

Using the Netstat command

You can use the Netstat command to verify your TCP/IP configuration. The information provided in the output from the Netstat command should be checked against the values in your configuration data sets for the TCP/IP stack. See the PROFILE DD statement in the TCP/IP started task procedure for the name of the configuration data sets.

Netstat can be invoked by using the TSO NETSTAT command, the z/OS UNIX shell **netstat/onetstat** command, or the console DISPLAY TCPIP,,NETSTAT command.

The following Netstat commands can be used to verify the state of those network resources that affect connectivity:

- "Netstat HOME/-h"
- "Netstat DEVLINKS/-d"
- "Netstat ROUTE/-r"
- "Netstat ARP/-R" on page 40
- "Netstat ND/-n" on page 40

For a complete description of the Netstat command and examples of Netstat output, see z/OS Communications Server: IP System Administrator's Commands.

Netstat HOME/-h

Use the Netstat HOME/-h command to verify the IP addresses defined for a TCP/IP stack, the names of the interfaces which are associated with the IP addresses, and the status of the IPv6 IP addresses. If any of the displayed information appears incorrect, check the HOME and INTERFACE statements in the PROFILE.TCPIP data set.

Netstat DEVLINKS/-d

Use the Netstat DEVLINKS/-d command to display the status and associated configuration values for a device and its defined interfaces, as coded in the PROFILE.TCPIP data set.

Netstat ROUTE/-r

The Netstat ROUTE/-r command displays the current routing tables for TCP/IP. To establish connectivity to a remote host, the remote host must also have a route back to the z/OS Communications Server.

The Netstat ROUTE/-r RSTAT command displays all of the static routes that are defined as replaceable.

The Netstat ROUTE/-r RADV command displays all of the IPv6 routes added based on information received in router advertisement messages.

The Netstat ROUTE/-r PR command displays all of the routes available in policy-based routing tables.

If there are any errors in the policy-based routing tables, check policy agent startup and configuration files for probable errors.

- Ensure that no error messages were generated during processing of either the initial profile or any subsequent VARY TCPIP,,OBEYFILE commands. (For information about the VARY TCPIP,,OBEYFILE command, see z/OS Communications Server: IP System Administrator's Commands.)
- Check the PROFILE.TCPIP data set for the following:
  - Ensure that the HOME and INTERFACE statements have been coded correctly.
  - If static routing is provided by using the BEGINROUTES statement, ensure that following conditions are true:
    - Each route in the statement correlates to a valid interface name.
    - There are routes in the statement that correlate to the appropriate network and host addresses available on the network.

Netstat ARP/-R

Use the command Netstat ARP/-R to query the ARP cache for a given address. Use Netstat ARP/-R ALL to query an entire ARP cache table. Ensure Netstat ARP/-R displays an ARP entry for the remote hosts.

The ARP entry for the host on a remote network contains the IP address and the MAC address for the router.

To ensure the host has a route back to the z/OS Communications Server, review the routing tables on the remote host. The route back can be a host route or network route. Intermediate routers must also be configured correctly.

Netstat ND/-n

Use Netstat ND/-n to display the Neighbor Discovery entries.

Using the DISPLAY TCPIP,,OSAINFO command

You can use the DISPLAY TCPIP,,OSAINFO command to retrieve information for an active IPAQENET/IPAQENET6 interface or link directly from an OSA-Express feature. The report includes information about the OSA-Express feature and about the interface or link. Registered unicast and multicast addresses are displayed as are routing variables for QDIO inbound workload queueing. Some of these values can be compared to the Netstat DEvlinks/-d report to ensure z/OS Communications Server and the OSA-Express are using the same information.

For a complete description of the command and examples of the command output, see DISPLAY TCPIP,,OSAINFO in z/OS Communications Server: IP System Administrator's Commands.

Using the Traceroute command

Traceroute displays the route that a packet takes to reach the requested target. Traceroute starts at the first router and uses a series of UDP probe packets with increasing IP time-to-live (TTL) or hop count values to determine the sequence of routers that must be traversed to reach the target host. The Traceroute function can be invoked by either the TSO TRACERTE command or the z/OS UNIX shell **traceroute/otracert** command.

The packetSize option lets you increase the IP packet size to see how size affects the route that the Traceroute packet takes. It also shows the point of failure if a destination address cannot be reached.

If equal-cost multipath routes exist in the IP routing table for outbound IP traffic to reach a remote host, use the Traceroute SRCIP/-s option or the INTF/-i option to select a home IP address (for example, VIPA) for the source IP address and a routing interface with the attached equal-cost multipath route. Alternatively, for routing interfaces associated with an IPv6 link-local address, you can append the name of the routing interface as scope information to the IPv6 link-local address of the remote host. When running multiple TCP/IP stacks on the same MVS image, specify the TCP/-a parameter, with the scope, to indicate the stack to which the routing interface is configured. Whenever applicable, use one of these options to test connectivity. For more information about using scope, see the information about support for scope in z/OS Communications Server: IPv6 Network and Application Design Guide.

For the complete syntax of the TSO TRACERTE and z/OS UNIX **traceroute/otracert** command and examples of command output, see the z/OS Communications Server: IP System Administrator's Commands.

Using SNMP remote Ping command

Use the SNMP remote Ping command to determine the response time between two remote hosts. For example, from Host A, you can determine the response time (Ping) between Hosts B and C, assuming the SNMP agent and TCP/IP subagent are running on Host B. See the z/OS Communications Server: IP System Administrator's Commands.

Documentation for the IBM Support Center

In most cases, persistent error conditions indicate an installation or configuration problem. Contact the local IBM branch office for installation assistance.

If a software defect is suspected, collect the following information before contacting the IBM Support Center:

- PROFILE.TCPIP
- TCPIP.DATA
- Output from Netstat commands. If using policy-based routing, collect Netstat ROUTE/-r output for all possible route tables involved in the failed routing.
- Output from Ping traces
- If using policy-based routing, output from **pasearch** commands
- Network diagram or layout
- Error messages received. See z/OS Communications Server: IP Messages Volume 4 (EZZ, SNM) for information about messages.

- Component traces, see Chapter 5, “TCP/IP services traces and IPCS support,” on page 45
- If using dynamic routing protocols for IP route table management, see the following information:
  - Chapter 31, “Diagnosing OMPROUTE problems,” on page 759
  - Chapter 32, “Diagnosing NCPROUTE problems,” on page 793

Part 2. Traces and control blocks

© Copyright IBM Corp. 2000, 2015

Chapter 5. TCP/IP services traces and IPCS support

This topic describes selected procedures for TCP/IP Services component trace, packet trace, and Socket API trace. The following subtopics are included:

- "Component trace"
- "Event trace (SYSTCPIP) for TCP/IP stacks and Telnet" on page 58
- "Packet trace (SYSTCPDA) for TCP/IP stacks" on page 93
- "Data trace (SYSTCPDA) for TCP/IP stacks" on page 146
- "Intrusion Detection Services trace (SYSTCPIS)" on page 149
- "OSAENTA trace (SYSTCPOT)" on page 166
- "Network security services (NSS) server trace (SYSTCPNS)" on page 170
- "Defense Manager daemon (DMD) trace (SYSTCPDM)" on page 170
- "OMPROUTE trace (SYSTCPRT)" on page 170
- "RESOLVER trace (SYSTCPRE)" on page 170
- "Configuration profile trace" on page 171

The TN3270E Telnet server uses a subset of the TCP/IP Services component trace. Specify the started procedure name of Telnet instead of TCP/IP to control component tracing in the Telnet address space.

Component trace

You typically use component trace when recreating a problem.

Component trace performs the following functions:

- Captures trace requests.
- Adds trace records to an internal buffer.
- Writes the internal buffer to an external writer, if requested.
- Formats the trace records using the Interactive Problem Control System (IPCS) subcommand CTRACE.
- Provides a descriptor at the beginning of a trace record that specifies the address and length of each data area. Each data area in the trace record is dumped separately.
- Provides an optional identifier for the connection (UDP, TCP, and so on) as part of each record.

Tip: Trace data can contain user IDs, passwords, and other sensitive information. The trace data files should be protected to prevent disclosure. As an example, packet trace of the FTP port 21 used to control FTP sessions contains user IDs and passwords in the CLEAR. However, a customer can use Secure Socket Layer for FTP and for TELNET. The Packet Trace (V TCPIP,,PKTTRACE) command can be RACF® protected.

For detailed information, see the following information:

- z/OS MVS Diagnosis: Tools and Service Aids for information about component trace procedures.
- z/OS MVS Initialization and Tuning Reference for information about the component trace SYS1.PARMLIB member.

© Copyright IBM Corp. 2000, 2015

- z/OS MVS System Commands for information about commands.
- z/OS MVS Programming: Authorized Assembler Services Guide for procedures and return codes for component trace macros.
- z/OS MVS IPCS Commands for information about IPCS commands.
- z/OS MVS IPCS User's Guide for information about using IPCS.

Modifying options with the TRACE CT command

After initialization, you must use the TRACE CT command to change the component trace options. Modifying options with the TRACE CT command can be done with or without the PARMLIB member. The component trace buffer size can be changed for the SYSTCPDA, SYSTCPIP, SYSTCPIS, and SYSTCPOT components.

Modifying with the PARMLIB member

Because TCP/IP, OMPROUTE, RESOLVER, IKE daemon, NSS server, DMD, and the trace command are accessing the PARMLIB data sets, they need to be authorized for read access to these data sets by RACF or another security product.

To change component trace options using a PARMLIB member, create a SYS1.PARMLIB member and specify the component member on the PARM= keyword of the TRACE CT command.

Use the following syntax:

```
TRACE CT,ON,COMP=component_name,SUB=(procedure_jobname)
,PARM=parmlib_member
```

Following are descriptions of the parameters:

COMP

Indicates the component name:

SYSTCPDA

TCP/IP packet trace. There is no parmlib member. Options are specified by the VARY TCPIP,,PKTTRACE command (see “Packet trace (SYSTCPDA) for TCP/IP stacks” on page 93).

SYSTCPDM

Defense Manager daemon, parmlib = CTIDMD00 (see “TCP/IP services component trace for the Defense Manager daemon” on page 728).

SYSTCPIK

IKE daemon, parmlib = CTIIKE00 (see “TCP/IP services component trace for the IKE daemon” on page 351).

SYSTCPIP

TCP/IP event trace, parmlib = CTIEZB*xx*, where xx is any 2 alphanumeric characters (see “Event trace (SYSTCPIP) for TCP/IP stacks and Telnet” on page 58).

SYSTCPIS

TCP/IP intrusion detection service, parmlib = CTIIDS*xx* (see “Intrusion Detection Services trace (SYSTCPIS)” on page 149).

SYSTCPNS

Network security services server, parmlib = CTINSS00 (see “TCP/IP services component trace for the network security services (NSS) server” on page 377).

SYSTCPOT
TCPIP OSA-Express Network Traffic Analyzer (OSAENTA) trace. TCP/IP event trace, parmlib = CTINTA00, (see "OSAENTA trace (SYSTCPOT)" on page 166). An alternate CTINTA00 member cannot be specified on the EXEC statement of the TCPIP procedure. CTINTA00 will always be used when starting TCPIP. Only an alternate buffer size or external writer procedure can be specified. All options are provided by the OSAENTA command.

SYSTCPRE
Resolver, parmlib = CTIRES*xx* (see Chapter 39, "Diagnosing resolver problems," on page 869).

SYSTCPRT
OMPROUTE, parmlib = CTIORA00 (see "TCP/IP services component trace for OMPROUTE" on page 786).

Tip: An optional suffix, CTIORAxx, is also available.

SUB
Indicates the started procedure name for TCP/IP, the OMPROUTE application, the RESOLVER, the IKE daemon started task name, the network security services (NSS) server started task name, the Defense Manager daemon (DMD) started task name, or the Telnet started task name for which the trace is run. If you use the S *procname,jobname* method to start TCP/IP, OMPROUTE, IKE daemon, network security services (NSS) server, DMD, or Telnet, you must specify the same value for the SUB parameter that is specified for the *jobname* value. There can be as many as eight TCP/IP sessions and eight Telnet sessions active in one system.

Restrictions:
- Only one OMPROUTE application can be active on each TCP/IP stack.
- Only one RESOLVER application can be active with each operating system.
- Only one IKE daemon application can be active with each operating system.
- Only one network security services (NSS) server application can be active with each operating system.
- Only one Defense Manager daemon (DMD) application can be active with each operating system.

PARM
Identifies the PARMLIB member that contains the trace options (see COMP). All options can be respecified. However, the buffer size cannot be changed if OMPROUTE, the IKE daemon, the NSS server, the DMD, or the RESOLVER are running. If a different size is required, you must stop OMPROUTE, IKE daemon, network security services server, DMD, or the RESOLVER, and then restart it after modifying the PARMLIB member.

If the incorrect parmlib member is specified, one of the following messages might be issued:

- An incorrect CTIEZB*xx* member is specified on the TRACE CT,ON command:

```
IEE538I CTIEZBxx MEMBER NOT FOUND IN SYS1.PARMLIB
ITT010I COMPONENT TRACE PROCESSING FAILED FOR PARMLIB MEMBER=CTIEZBxx:
PARMLIB MEMBER NOT FOUND.
```

- An incorrect CTIEZB*xx* member is specified on the CTRACE() keyword of the EXEC statement of the TCP/IP started procedure:

```
IEE538I CTIEZBxx MEMBER NOT FOUND IN SYS1.PARMLIB
```

- An incorrect CTIORA*xx* member is specified on the TRACE CT,ON command:

  ```
  IEE5381 CTIORAxx MEMBER NOT FOUND in SYS1.PARMLIB
  ITT01011 COMPONENT TRACE PROCESSING FAILED FOR PARMLIB MEMBER=CTIORAxx:
  PARMLIB MEMBER NOT FOUND
  ```

- An incorrect CTINTA00 member is specified on the TRACE CT,ON command:

  ```
  IEE5381 CTINTA00 MEMBER NOT FOUND in SYS1.PARMLIB
  ITT01011 COMPONENT TRACE PROCESSING FAILED FOR PARMLIB MEMBER=CTINTA00:
  PARMLIB MEMBER NOT FOUND
  ```

Modifying without the PARMLIB member

To change component trace options without using a PARMLIB member, issue the TRACE CT command without the PARM= parameter and specify the options on the reply. Though the SYSTCPDA component for packet or data trace does not have a parmlib member, SYSTCPDA can be used on the trace command without the PARMLIB member.

Use the following syntax:

```
TRACE CT,ON,COMP=component_name,SUB=(procedure_jobname)
```

After issuing the TRACE CT command, you are prompted to specify the trace options. Respond using the following syntax:

```
Reply nn
[,ASID=(asid-list)]
[,JOBNAME=(jobname-list)]
[,OPTIONS=(name[name]...)]
[,WTR={membername|DISCONNECT}]
[,CONT|END]
```

Restriction: ASID and JOBNAME are not valid for OMPROUTE.

Reply *nn*
: Specifies the identification number (in the range 0-9999) in the prompting message. For example, if the response is

  ```
  06 ITT066A SPECIFY OPERAND(S) FOR TRACE CT COMMAND
  ```

  You might reply

  ```
  r 06,WTR=PTTCP,END
  ```

ASID
: The ASID (address space identifiers) of the client whose TCP/IP requests are to be traced.

JOBNAME
: The job name of the client whose TCP/IP requests are to be traced. The job name might be:

  - The job name associated with a client application.
  - The SNA LU associated with a TELNET session.

    Restriction: Do not use the JOBNAME parameter with the TELNET CTRACE option.
  - The FTP user ID associated with an FTP data connection.

OPTIONS
: Options valid for use with SYSTCPIP are listed in this topic; options valid for use with OMPROUTE are listed in Chapter 31, “Diagnosing OMPROUTE problems,” on page 759; and options for SYSTCPRE (the Resolver component) are listed in Chapter 39, “Diagnosing resolver problems,” on page 869.

Options valid for use with IKE daemon are listed in Chapter 9, “Diagnosing IKE daemon problems,” on page 333

Options valid for use with the network security services (NSS) server are listed in Chapter 10, “Diagnosing network security services (NSS) server problems,” on page 357.

Options valid for use with the Defense manager daemon (DMD) are listed in Chapter 29, “Diagnosing Defense Manager daemon problems,” on page 721.

membername
The member containing the source JCL that invokes the external writer. The *membername* in the WTR parameter must match the *membername* in a previous TRACE CT,WTRSTART command. (See “Steps for obtaining component trace data with an external writer” on page 51.)

WTR=DISCONNECT
Disconnects the component trace external writer and the trace. You must also specify a TRACE CT,WTRSTART or TRACE CT,WTRSTOP command to start or stop the writer.

CONT or END
CONT specifies that the reply continues on another line. Specify END to complete the response.

Displaying component trace status

To display information about the status of the component trace, issue the following command:

```
DISPLAY TRACE,COMP=component_name,SUB=(procedure_jobname)
```

See COMP for more information about *component_name*.

This command displays information about the status of the component trace for one procedure. To display information about the status of the component trace for all active procedures, issue the following command:

```
DISPLAY TRACE,COMP=component_name,SUBLEVEL,N=8
```

For the TCP/IP CTRACE components, do not be misled by the line in the middle of the display showing the MODE is OFF. This part of the display always says the MODE is OFF because TCP/IP uses the subtrace for all tracing. The subtrace for TCPCS2 indicates the actual state of the trace. In the example shown in 4 on page 52, the trace is active (MODE is ON) with an internal buffer size of 16 M, tracing all ASIDs and all JOBNAMES, using MINIMUM options, and using the external writer PTTCP. Another version of the DISPLAY TRACE command D TRACE,COMP=*component_name*,SUBLEVEL,N=8 shows all subtraces for the component.

Modifying the trace buffer size:
To modify the amount of trace buffer in use for the SYSTCPIP, SYSTCPDA, SYSTCPIS and SYSTCPOT traces use the following command:

```
TRACE CT,nnnnM,COMP=component_name,SUB=(procedure_jobname)
```

where *nnnnM* is the new buffer size in mega bytes. The buffer size is subject to the minimum and maximum buffer size established for each component.

See COMP for more information about *component_name*.

Stopping a component trace

To stop current tracing, issue the following TRACE CT command:

```
TRACE CT,OFF,COMP=component_name,SUB=(procedure_jobname)
```

See COMP for more information about *component_name*.

With the TRACE,CT,OFF command, TCP/IP discontinues recording of all trace data.

```
TRACE CT,OFF,COMP=SYSTCPIP,SUB=(procedure_jobname)
```

Obtaining component trace data with a dump

You can request a dump to obtain component trace data for:

- TCP/IP stack
- OMPROUTE
- Resolver
- TELNET
- IKE daemon
- Network security services (NSS) server
- Defense Manager daemon (DMD)

TCP/IP stack:
If an abend occurs in the TCP/IP address space or in a user's address space, TCP/IP recovery dumps the home ASID, primary ASID, and the secondary ASID. This also contains the trace data for the SYSTCPIP, SYSTCPDA, SYSTCPIS and SYSTCPOT components.

To view the trace records for a problem where no abend has occurred, use the DUMP command. The following example illustrates a DUMP command:

```
DUMP COMM=(your dump title here)
R n,JOBNAME=(tcpipprocname),CONT
R n,SDATA=(NUC,CSA,LSQA,PSA,RGN,SQA,TRT),END
```

Figure 7. Example of DUMP command for TCP/IP stack

To generate a meaningful dump, specify (at a minimum):

- *CSA*
- *LSQA*
- *RGN*
- *SQA*

OMPROUTE:
To obtain a dump of the OMPROUTE address space (which contains the trace table), use the DUMP command, as shown in the following example:

```
DUMP COMM=(enter your dump title here)
R n,JOBNAME=omproute_started_task_name,SDATA=(CSA,RGN,ALLPSA,SQA,SUM,TRT,ALLNUC),END
```

Figure 8. Example of DUMP command for OMPROUTE

RESOLVER:
To obtain a dump of the RESOLVER, use the DUMP command, as shown in the following example:

```
DUMP COMM=(enter your dump title here)
R n,JOBNAME=resolver_started_task_name,SDATA=(CSA,RGN,ALLPSA,SQA,SUM,TRT,ALLNUC),END
```

Figure 9. Example of DUMP command for RESOLVER

TELNET:
To obtain a dump of TELNET, use the DUMP command, as shown in the following example:

```
DUMP COMM=(enter your dump title here)
R n,JOBNAME=telnet_started_task_name,SDATA=(CSA,RGN,ALLPSA,SQA,SUM,TRT,ALLNUC),END
```

Figure 10. Example of DUMP command for TELNET

Steps for obtaining component trace data with an external writer

You can use an external writer to obtain component trace data for TCP/IP stacks, packet trace, OMPROUTE, and Telnet.

Procedure

Perform the following steps to obtain component trace data with an external writer:

1. Enter the appropriate writer procedure in SYS1.PROCLIB, as shown in the following example. Use a separate external writer for each CTRACE component. You can have multiple procedures writing to as many as 16 TRCOUT files either on disk or tape.

   ```
   //PTTCP  PROC
   //* REFER:  SYS1.PROCLIB(PTTCP)
   //* COMPID: OPER
   //* DOC:    THIS PROCEDURE IS THE IPCS CTRACE1 EXTERNAL WRITER PROCEDURE.
   //*         USED BY TCP/IP   .
   //*
   //IEFPROC  EXEC  PGM=ITTTRCWR,REGION=0K,TIME=1440
   //* TIME=1440 to prevent S322 abends
   //TRCOUT01 DD DSNAME=MEGA.IPCS.CTRACE1,UNIT=SYSDA,
   //             VOL=SER=STORGE,
   //             SPACE=(4096,(100,10),,CONTIG),DISP=(NEW,CATLG),
   // DCB=(DSORG=PS)
   //
   ```

   Restrictions:

   - Do not specify the DCB parameters RECFM, LRECL, and BLKSIZE. The external writer defaults to an optimal blocking factor.
   - Do not specify the RLSE subparameter of the SPACE parameter. The output data sets are opened and closed more than once when GTF runs.

2. Start the external writer using the following command:

   ```
   TRACE CT,WTRSTART=procedure_name,WRAP
   ```

3. Turn on the trace and connect the external writer to the component either by specifying the external writer name in the PARMLIB member, or by specifying the external writer name in the TRACE command. When starting TCP/IP, because the SYSTCPDA component has no PARMLIB member, the PARMLIB option is not applicable for SYSTCPDA. For example, `TRACE CT,ON,COMP=SYSTCPDA,SUB=(TCPCS),PARM=CTIEZBDA` is a valid command. The PARMLIB member can specify a new buffer size or the name or a writer. To turn the trace on and connect the external writer to the component using a PARMLIB member, add the following TRACEOPTS option to the PARMLIB member:

   ```
   WTR(xxx)
   ```

where *xxx* is the procedure name of the external writer. Then use this PARMLIB member when starting the program (TCP/IP, OMPROUTE, TELNET, or the Resolver) or if the program is already executing, issue the following command:

```
TRACE CT,ON,COMP=component_name,SUB=(procedure_name),PARM=parmlib_member
```

To turn on the trace and connect the external writer without using the PARMLIB member, enter the following command:

```
TRACE CT,ON,COMP=component_name,SUB=(procedure_name)
```

When the system responds, enter the following command:

```
R n,WTR=procedure_name,END
```

where *n* is the response number issued by the system. You can add options to the response. The options vary for each component name. See "Formatting component traces" on page 53 for references to the component options.

4. Use the DISPLAY command to check the external writer status. Include a sublevel.

```
D TRACE,COMP=SYSTCPDA,SUB=(TCPCS2)
IEE843I 11.33.06  TRACE DISPLAY 099
        SYSTEM STATUS INFORMATION
 ST=(ON,0064K,00064K) AS=ON  BR=OFF EX=ON  MT=(ON,064K)
  TRACENAME
  =========
  SYSTCPDA
                       MODE BUFFER HEAD SUBS
                       ====================
                       OFF         HEAD    2
     NO HEAD OPTIONS
  SUBTRACE             MODE BUFFER HEAD SUBS
 ------------------------------------------------------------
  TCPCS2               ON   0016M
     ASIDS      *NONE*
     JOBNAMES   *NONE*
     OPTIONS    MINIMUM
     WRITER     PTTCP
```

Tip: The external writer is active for packet and data.

5. Turn off the trace or disconnect the external writer. The following two commands disconnect from the external writer, while leaving the trace running internally.

```
TRACE CT,ON,COMP=component_name,SUB=(procedure_jobname)
```

When the system responds, enter the second command:

```
R nn,WTR=DISCONNECT,END
```

6. Stop the external writer using the following command:

```
TRACE CT,WTRSTOP=procedure_name
```

Tips for using component trace external writer

Consider the following when using the component trace external writer:

- Do not use the same writer to trace more than one TCP/IP stack, TELNET, or OMPROUTE application. If you need to trace multiple stacks or applications, use separate writers.
- If your external writer fills up and the wrap option is on, the writer overwrites itself. If the nowrap option is on, the writer stops.
- Use REGION=0K on the trace writer procedure EXEC statement. This helps ensure that there is enough virtual memory for trace buffers.
- Use TIME=1440 on the EXEC statement. This prevents S322 abends.

- Use CONTIG on the disk space allocation of the trace data when using the WRAP option. For example: SPACE=(1024,(4096,100),,CONTIG). This ensures that the space for the trace data set is available.
- Do not specify DCB parameters for trace data sets. The writer optimizes the logical record length and block size for new trace data sets.
- Ensure that the dispatching priority of the writer is equal to or greater than the application that is being traced.
- The CTRACE or packet trace formatter messages ITT10016I, PTRPT14I or PTRPT15I indicate lost buffers. This normally occurs when packet tracing is set up with no capture filters, causing all protocols for all IP addresses to be captured. The result is that some trace buffers are unable to be processed and are flagged as lost. The reason behind these lost records is the amount of data that is being traced over-running the CTRACE writer. This trace data is arriving at the CTRACE writer faster than it can be written out. Setting filters such as IPADDR or PORTNUM will help eliminate lost records.

Using a VSAM linear data set:

Using a VSAM linear data set for output trace data provides better performance than using a sequential data set.

Procedure

Perform the following steps:

1. Define a VSAM Linear data set.

```
//DEFINE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
 DELETE +
     (hlq.CTRACE.LINEAR)     +
   CLUSTER
 DEFINE CLUSTER( +
 NAME(hlq.CTRACE.LINEAR)     +
 LINEAR                      +
 MEGABYTES(10)               +
 VOLUME(CPDLB0)              +
 CONTROLINTERVALSIZE(32768)  +
 )                           +
 DATA(                       +
 NAME(hlq.CTRACE.DATA)       +
  )
LISTCAT ENT(hlq.CTRACE.LINEAR) ALL
```

2. Update the CTRACE writer procedure:

```
//IEFPROC  EXEC PGM=ITTTRCWR
//TRCOUT01 DD DSNAME=hlq.CTRACE.LINEAR,DISP=SHR
//SYSPRINT DD SYSOUT=*
```

3. Issue the COPYTRC command. The VSAM data set must be copied with COPYTRC to a sequential data set before being sent to IBM Service.

Formatting component traces

You can format component trace records using IPCS panels or a combination of IPCS panels and the CTRACE command, either from a dump or from external-writer files. The code for the component trace record formatter can be found in the SYS1.MIGLIB data set. This data set should be added as a concatenation to the STEPLIB data set. For details, see z/OS MVS IPCS Commands and z/OS MVS IPCS User's Guide.

Steps for formatting component traces using IPCS panels:

Format component traces by using only IPCS panels.

Procedure

Follow these steps:

1. Log on to TSO.
2. Access IPCS.
3. Select option 2 (ANALYSIS) from the option list.
4. Select option 7 (TRACES) from the option list.
5. Select option 1 (CTRACE) from the option list.
6. Select option D (Display) from the option list.

Results

You know you are done when the CTRACE DISPLAY PARAMETERS screen is displayed (Figure 11), as shown below.

```
ITTPC503----------CTRACE DISPLAY PARAMETERS-------------------------

System      =======>           (System name or blank)
Component   =======>           (Component name (required))
Subnames    =======>

GMT/LOCAL   =======>            (Greenwich Mean Time or Local; GMT is default)
Start time  =======>            (mm/dd/yy,hh:mm:ss.dddddd)
Stop time   =======>
Limit       =======>      Exception  =======>
Report type =======> FULL       (SHort, SUmmary, Full, Tally)
User exit   =======>            (Exit program name)
Override source =======>
Options     =======>

To enter/verify required values, type any character
Entry IDS =======>     Jobnames =======>     ASIDs =======>     OPTIONS =======>     SUBS =======>

CTRACE COMP(xx) FULL

COMMAND =======>
     F1=Help  F2=Split  F3=End    F4=RETURN    F5=RFIND    F6=MORE    F7=UP
     F8=DOWN  F9=Swap  F10=LEFT  F11=RIGHT    F12=CURSOR
```

Figure 11. IPCS CTRACE

Enter the component name in the COMPONENT field and as the value in COMP(*xx*). For descriptions of options, see the following topics:

- SYSTCPDA, see COMP.
- SYSTCPDM, see "TCP/IP services component trace for the Defense Manager daemon" on page 728.
- SYSTCPIK, see "TCP/IP services component trace for the IKE daemon" on page 351.
- SYSTCPNS, see "TCP/IP services component trace for the network security services (NSS) server" on page 377.
- SYSTCPIP, see COMP.
- SYSTCPIS, see COMP.
- SYSTCPOT, see "OSAENTA trace (SYSTCPOT)" on page 166.
- SYSTCPRE, see Chapter 39, "Diagnosing resolver problems," on page 869.

- SYSTCPRT, see "TCP/IP services component trace for OMPROUTE" on page 786.

Steps for using the CTRACE command:

Format component traces by using the CTRACE command.

Procedure

Perform the following steps:

1. Log on to TSO.
2. Access IPCS.
3. Select option 6 (COMMAND) from the option list.
4. Enter a CTRACE command and options on the IPCS command line.

Syntax:
Following is the syntax of the IPCS CTRACE command:

CTRACE syntax

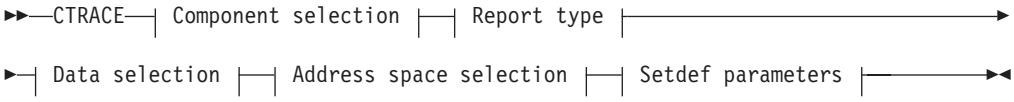

Component Selection:

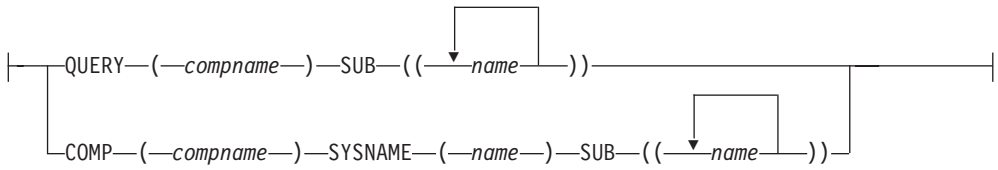

Report Type:

Data Selection:

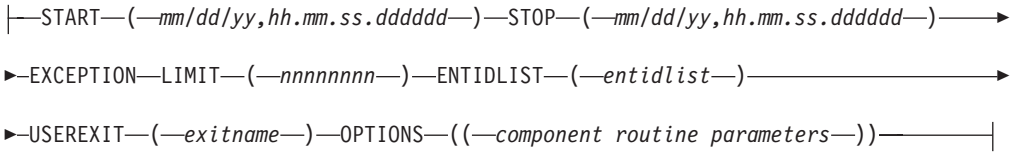

Address Space Selection:

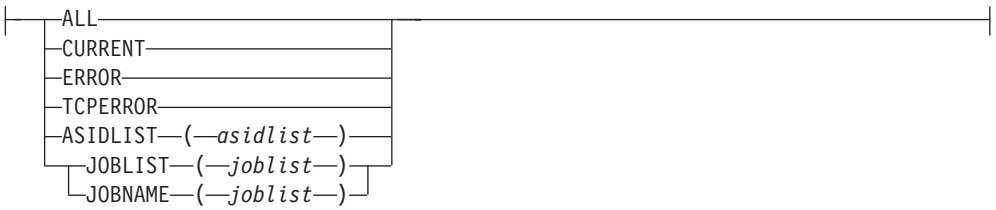

Setdef Parameters:

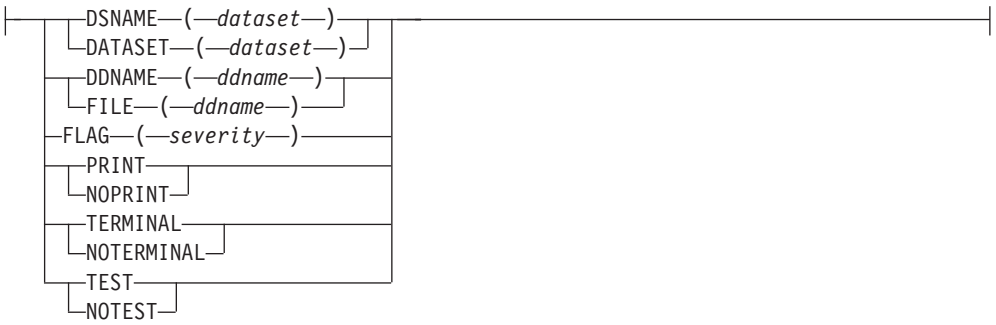

Parameters:
See z/OS MVS IPCS Commands for details on the CTRACE parameters.

Keywords:
You can use the following CTRACE keywords with TCP/IP component trace formats:

JOBLIST, JOBNAME
Use the JOBLIST and JOBNAME keywords to select packet trace records with a matching link name. However, only the first 8 characters of the link name are matched and no asterisks are accepted in the job name. Also, use them to match the job name in data trace records.

ASIDLIST
Use the ASIDLIST to select trace records only for a particular address space.

GMT
The time stamps are converted to GMT time.

LOCAL
The time stamps are converted to LOCAL time.

SHORT
If the OPTIONS string does not specify any reports, then format the trace records. Equivalent to the FORMAT option.

FULL
If the OPTIONS string does not specify any reports, then format and dump the trace records. Equivalent to the FORMAT and DUMP options.

SUMMARY
If the OPTIONS string does not specify any reports, then create a one line summary for each trace record. Equivalent to the SUMMARY option.

TALLY
If the OPTIONS string does not specify any reports, then count the trace records.

START and STOP
These keywords limit the trace records that are seen by the formatter. The STOP keyword determines the time when records are no longer seen by the packet trace report formatter.

Rule: CTRACE always uses the time the trace record was moved to the buffer for START and STOP times.

LIMIT
Determines the number of records the formatter is allowed to process.

USEREXIT
The CTRACE USEREXIT is called for TCP/IP formatter, except for the packet trace formatters. Therefore, the packet trace formatter calls the CTRACE USEREXIT before testing the records with the filtering criteria. If it returns a nonzero return code, then the record is skipped. The USEREXIT can also be used in the OPTIONS string. It is called after the record has met all the filtering criteria in the OPTIONS string. For details, see "Formatting packet traces using IPCS" on page 96.

Examples of formatting component traces:
The following example shows the error message when the specified address space is not available in the dump.

```
CTRACE QUERY(SYSTCPIP) SUB((TCPSVT1)) FULL LOCAL
COMPONENT TRACE QUERY SUMMARY

ITT10003I There are no trace buffers in the dump for COMP(SYSTCPIP)SUB((TCPSVT1))
```

The following example shows the results when the CTRACE QUERY command is issued for a dump when the address space is available.

```
CTRACE QUERY(SYSTCPIP) SUB((TCPSVT2)) FULL LOCAL
COMPONENT TRACE QUERY SUMMARY

 COMP(SYSTCPIP)SUBNAME((TCPSVT2))

               START = MT  02/21/2001 15:25:49.432  LOCAL
               STOP  = 180   02/21/2001 15:51:16.8
             Buffer size: 0050M

             OPTIONS: CONFIG,CSOCKET,FIREWALL,IOCTL,MESSAGE,OETCP,OPCMDS,
                      OPMSGS,PASAPI,PING,SOCKAPI,TN,UDP,XCF,CLAW,INT
             ERNET,LCS,VTAM,VTAMDATA

             OPTIONS: MINIMUM
```

Tip: The first option is the relevant one (ignore the second options list). The buffer size and options list are displayed only for a dump data set, not an external writer data set.

Formatting component traces using a batch job:
A component trace can also be formatted through the use of a batch job. The following is an example of JCL for a batch job:

```
//jobname   DD (accounting),pgmname,CLASS=A,MSGCLASS=A
//DUMP    EXEC PGM=IKJEFT01
//STEPLIB DD DISP=SHR,DSN=SYS1.MIGLIB
//SYSPRINT  DD SYSOUT=*
//SYSUDUMP  DD SYSOUT=*
//SYSTSPRT  DD SYSOUT=*
//PRINTER   DD SYSOUT=*
//SYSPROC   DD DISP=SHR,DSN=SYS1.CLIST
//          DD DISP=SHR,DSN=SYS1.SBLSCLI0
```

```
//IPCSPARM  DD DISP=SHR,DSN=SYS1.PARMLIB
//          DD DISP=SHR,DSN=CPAC.PARMLIB
//          DD DISP=SHR,DSN=SYS1.IBM.PARMLIB
//IPCSPRNT  DD SYSOUT=*
//IPCSTOC   DD SYSOUT=*
//IPCSDDIR  DD DISP=SHR,DSN=userid.IPCS.DMPDIR
//SYSTSIN   DD *
 IPCS NOPARM
 SETDEF DA('ctrace.dataset')
 CTRACE COMP(SYSTCPIP) SUBNAME((tcpiprocname)) OPTIONS((systcpip_options_string)) +
  FULL LOCAL
 END
/*

Note: IPCSPARM DD should be modified as follows:
//IPCSPARM DD DISP=SHR,DSN=SYS1.PARMLIB
//         DD DISP=SHR,DSN=CPAC.PARMLIB
//         DD DISP=SHR,DSN=SYS1.IBM.PARMLIB

These concatenations will be used to locate the BLSCECT member that
is required by IPCS
```

IKE daemon trace (SYSTCPIK)

TCP/IP Services component trace is also available for use with the IKE daemon. See "TCP/IP services component trace for the IKE daemon" on page 351.

Event trace (SYSTCPIP) for TCP/IP stacks and Telnet

The TN3270E Telnet server running as its own procedure also uses the SYSTCPIP event trace.

Restrictions: All discussion that follows where TCP/IP is used as an example also pertains to the TN3270E Telnet server with the following exceptions:

- The TN3270E Telnet server does not use 64-bit common storage for trace collection, it uses its own private storage.
- A subset of trace commands are used by Telnet. A default parmlib member, CTIEZBTN, is provided that indicates all trace options available. The default parmlib member can be overridden in the same manner as the TCP/IP parmlib can be overridden.
- A subset of IPCS commands are used by Telnet.

Event trace for TCP/IP stacks traces individual TCP/IP components (such as STORAGE, INTERNET, and so forth) and writes the information either to a data set (using an external writer), or internally to 64-bit common storage. To aid in first failure data capture, a minimal component trace is always started during TCP/IP initialization if you use the TCP/IP Component Trace SYS1.PARMLIB member, CTIEZB*xx*.You can select trace records at run time by any of the following methods:

- JOBNAME
- Address space identifiers (ASID)
- Trace option
- IP address
- Port number
- Event identifier

Restriction: If using the TELNET options, do not specify the JOBNAME parm when starting CTRACE.

Specifying trace options

You can specify component trace options at TCP/IP initialization or after TCP/IP has initialized.

Specifying trace options at initialization

To start TCP/IP with a specific trace member, use the following command:

S *tcpip_procedure_name*,PARMS=CTRACE(CTIEZB*xx*)

where CTIEZB*xx* is the component trace SYS1.PARMLIB member.

You can create this member yourself, or you can update the default SYS1.PARMLIB member, CTIEZB00. For a description of trace options available in the CTIEZB00, see Table 11 on page 61.

Tip: Besides specifying the TCP/IP traces that you want, you can also change the component trace buffer size.

You can use IBM Health Checker for z/OS to check whether TCP/IP Event Trace (SYSTCPIP) is active with options other than the default options (MINIMUM, INIT, OPCMDS, or OPMSGS). For more details about IBM Health Checker for z/OS, see Appendix C, "IBM Health Checker for z/OS," on page 999.

```
/*********************************************************************/
/*                                                                   */
/*  IBM Communications Server for z/OS                               */
/*  SMP/E Distribution Name: CTIEZB00                                */
/*                                                                   */
/*  MEMBER:  CTIEZB00                                                */
/*                                                                   */
/*                                                                   */
/* Copyright:    Licensed Materials - Property of IBM                */
/*                                                                   */
/*               5694-A01                                            */
/*                                                                   */
/*               Copyright IBM Corp. 1996, 2011.                     */
/*                                                                   */
/*     STATUS =  CSV1R13                                             */
/*                                                                   */
/*  DESCRIPTION = This parmlib member causes component trace for     */
/*                the TCP/IP product to be initialized with a        */
/*                trace buffer size of 8 megabytes.                  */
/*                                                                   */
/*                This parmlib member only lists those TRACEOPTS     */
/*                value specific to TCP/IP.  For a complete list     */
/*                of TRACEOPTS keywords and their values see         */
/*                z/OS MVS INITIALIZATION AND TUNING REFERENCE.      */
/*                                                                   */
/* $MAC(CTIEZB00) PROD(TCPIP): Component Trace SYS1.PARMLIB member   */
/*                                                                   */
/*********************************************************************/
 TRACEOPTS
 /* ---------------------------------------------------------------- */
 /*   ON OR OFF: PICK 1                                              */
 /* ---------------------------------------------------------------- */
           ON
 /*        OFF                                                       */
 /* ---------------------------------------------------------------- */
 /*   BUFSIZE: A VALUE IN RANGE 1M TO 1024M                          */
 /* ---------------------------------------------------------------- */
           BUFSIZE(8M)
```

```
/*        JOBNAME(jobname1,...)                                        */
/*        ASID(Asid1,...)                                              */
/*        WTR(wtr_procedure)                                           */
/* ------------------------------------------------------------------- */
/*   Note, the following groups of trace options are supported:        */
/*                                                                     */
/*   ALL     = All options except MISC, PFSMIN, ROUTE, SERIAL,         */
/*             SOCKAPI, STORAGE, TCPMIN, and TIMER                     */
/*   CSOCKET = PFS + SOCKET                                            */
/*   DLC     = CLAW + INTERNET + LCS + VTAM + VTAMDATA                 */
/*   IN      = CONFIG + INIT + IOCTL + OPCMDS + OPMSGS                 */
/*   LATCH   = SERIAL                                                  */
/*   MINIMUM = INIT + OPCMDS + OPMSGS                                  */
/*   ALLMIN  = INIT + OPCMDS + OPMSGS + MINPFS + MINTCP                */
/*   OETCP   = ENGINE + PFS + QUEUE + TCP                              */
/*   OEUDP   = ENGINE + PFS + QUEUE + UDP                              */
/*   PING    = ARP + ICMP + RAW + ND                                   */
/*   RW      = ENGINE + PFS + QUEUE + RAW + SOCKET                     */
/*   SMTP    = ENGINE + IOCTL + PASAPI + PFS + QUEUE + SOCKET + TCP    */
/*   SYSTEM  = INIT + OPCMDS + OPMSGS + SERIAL + STORAGE + TIMER +     */
/*             WORKUNIT                                                */
/*   TC      = ENGINE + PFS + QUEUE + SOCKET + TCP                     */
/*   TN      = PFS + TCP + TELNET + TELNVTAM                           */
/*   UD      = ENGINE + PFS + QUEUE + SOCKET + UDP                     */
/*                                                                     */
/* ------------------------------------------------------------------- */
/*                                                                     */
/*   PFSMIN  = Reduced set of PFS trace data                           */
/*   TCPMIN  = Reduced set of TCP trace data                           */
/*   ALLMIN  = PFSMIN + TCPMIN                                         */
/*                                                                     */
/*                                                                     */
/*   NOTE:    The xxxMIN and the corresponding xxx options should      */
/*            not be active at the same time.  The will collect       */
/*            duplicate information.                                   */
/*                                                                     */
/*   OPTIONS: NAMES OF FUNCTIONS OR GROUPS TO BE TRACED:               */
/*                                                                     */
/* ------------------------------------------------------------------- */
/*        OPTIONS(             */
/*                 'ALL     '   */
/*               ,'ALLMIN  '   */
/*               ,'ACCESS  '   */
/*               ,'AFP     '   */
/*               ,'ARP     '   */
/*               ,'CLAW    '   */
/*               ,'CONFIG  '   */
/*               ,'CSOCKET '   */
/*               ,'DLC     '   */
/*               ,'EID(hhhhhhh,hhhhhhhh) '             */
/*               ,'ENGINE  '   */
/*               ,'FIREWALL'   */
/*               ,'ICMP    '   */
/*               ,'IN      '   */
/*               ,'INIT    '   */
/*               ,'INTERNET'   */
/*               ,'IOCTL   '   */
/*               ,'IPADDR(nnn.nnn.nnn.nnn/mmm.mmm.mmm, */
/*                        nnn.nnn.nnn.nnn/pp,          */
/*                        hhhh::hhhh/ppp)              */
/*               ,'IPSEC   '   */
/*               ,'LATCH   '   */
/*               ,'LCS     '   */
/*               ,'MESSAGE '   */
/*               ,'MINIMUM '   */
/*               ,'MISC    '   */
/*               ,'ND      '   */
```

```
/*              ,'NONE    '   */
/*              ,'OETCP   '   */
/*              ,'OEUDP   '   */
/*              ,'OPCMDS  '   */
/*              ,'OPMSGS  '   */
/*              ,'PASAPI  '   */
/*              ,'PFS     '   */
/*              ,'PFSMIN  '   */
/*              ,'PING    '   */
/*              ,'POLICY  '   */
/*              ,'PORT(ppppp,ooooo,rrrrr,ttttt) ' */
/*              ,'QUEUE   '   */
/*              ,'RAW     '   */
/*              ,'ROUTE   '   */
/*              ,'RW      '   */
/*              ,'SERIAL  '   */
/*              ,'SMTP    '   */
/*              ,'SNMP    '   */
/*              ,'SOCKAPI '   */
/*              ,'SOCKET  '   */
/*              ,'STORAGE '   */
/*              ,'SYSTEM  '   */
/*              ,'TC      '   */
/*              ,'TCP     '   */
/*              ,'TCPMIN  '   */
/*              ,'TELNET  '   */
/*              ,'TELNVTAM'   */
/*              ,'TIMER   '   */
/*              ,'TN      '   */
/*              ,'UD      '   */
/*              ,'UDP     '   */
/*              ,'VTAM    '   */
/*              ,'VTAMDATA'   */
/*              ,'WORKUNIT'   */
/*              ,'XCF     '   */
/*               )            */
```

Figure 12. SYS1.PARMLIB member CTIEZB00

A group activates multiple trace options. The group name identifies traces that should be activated for a specific problem area, and trace groups provide a way to collect trace data by problem type.

Table 11 describes the available trace options and groups.

Table 11. Trace options and groups

Trace Event	Description
ALL	All types of records except MISC, PFSMIN, ROUTE, SERIAL, STORAGE, TCPMIN, and TIMER. **Slow Performance:** Using this option slows performance considerably, so use with caution. Also available for the TN3270E Telnet server running in its own address space.

Table 11. Trace options and groups (continued)

Trace Event	Description
ALLMIN	Turns on the following trace options: • INIT • OPCMDS • OPMSGS • PFSMIN • TCPMIN
ACCESS	Trace creation, modification, and manipulation of the Network Access tree, along with results of all Network Access queries.
AFP	Turns on trace for fast response cache accelerator.
ARP	Shows address resolution protocol (ARP) cache management and ARP timer management. This option also shows all outbound and inbound ARP packets. **Tip:** The information provided differs depending on the type of device. **Guideline:** The ARP and ND options are aliases. If you turn one on, you turn on the other option, and if you turn one off, you turn off the other option. When formatting the trace, these options can be filtered separately.
CLAW	Shows all control flows for a CLAW device.
CONFIG	Turns on trace for configuration updates.
CSOCKET	Turns on the following trace options: • PFS • SOCKET
DLC	Turns on the following trace options: • CLAW • INTERNET • LCS • VTAM • VTAMDATA
EID(list)	Turns on trace by event identifier. The event identifiers are 8 hexadecimal digits. Up to 16 can be specified. Use only under the direction of IBM Support.
ENGINE	Turns on trace for stream head management. **Guideline:** The ENGINE and QUEUE options are aliases. If you turn one on, you turn on all related options, and if you turn one off, you turn off all related options. These alias options are only for recording the trace. When formatting the trace, these options can be filtered separately.
FIREWALL	Turns on trace for firewall events. **Tip:** Synonymous with IPSEC option.
ICMP	Turns on trace for the ICMP protocol.

Table 11. Trace options and groups (continued)

Trace Event	Description
IN	Turns on the following trace options: • CONFIG • INIT • IOCTL • OPCMDS • OPMSGS
INIT	Turns on trace for TCP/IP Initialization/Termination. **Note:** The INIT, OPCMDS, and OPMGS options are aliases. If you turn one on, you turn on all related options, and if you turn one off, you turn off all related options. These alias options are only for recording the trace. When formatting the trace, these options can be filtered separately. Also available for the TN3270E Telnet server running in its own address space.
INTERNET	Turns on trace for Internet Protocol layer. **Tip:** Using this option slows performance considerably, so use with caution.
IOCTL	Turns on trace for IOCTL processing.
IPADDR(list)	Turns on trace by IP address.
IPSEC	Turns on trace for IP security events. **Tip:** Synonymous with FIREWALL option.
LATCH	Turns on the following trace option: • SERIAL
LCS	Shows all control flows for an LCS device.
MESSAGE	Turns on trace for message triple management. **Tip:** Using this option slows performance considerably, so use with caution. Also available for the TN3270E Telnet server running in its own address space.
MINIMUM	Turns on the following trace options: • INIT • OPCMDS • OPMSGS
MISC	Turns on trace for miscellaneous TCP/IP internal diagnostic tests.
NONE	Turn off all traces but exception traces, which always stay on. Also available for the TN3270E Telnet server running in its own address space.
ND	Enable Neighbor Discovery trace option. **Guideline:** The ARP and ND options are aliases. If you turn one on, you turn on the other option, and if you turn one off, you turn off the other option. When formatting the trace, these options can be filtered separately.

Table 11. Trace options and groups (continued)

Trace Event	Description
OETCP	Turns on the following trace options: • ENGINE • PFS • QUEUE • TCP
OEUDP	Turns on the following trace options: • ENGINE • PFS • QUEUE • UDP
OPCMDS	Turns on traces of operator commands. **Guideline:** The INIT, OPCMDS, and OPMGS options are aliases. If you turn one on, you turn on all related options, and if you turn one off, you turn off all related options. These alias options are only for recording the trace. When formatting the trace, these options can be filtered separately.
OPMSGS	Turns on message trace for console messages. **Guideline:** The INIT, OPCMDS, and OPMGS options are aliases. If you turn one on, you turn on all related options, and if you turn one off, you turn off all related options. These alias options are only for recording the trace. When formatting the trace, these options can be filtered separately.
PASAPI	Turns on traces for transforms that handle Pascal APIs.
PFS	Turns on trace for the physical file system layer. **Tip:** The PFS and PFSMIN options should not be specified together; the PFS option gathers all the information that the PFSMIN option gathers.
PFSMIN	Turns on the minimum PFS trace option. **Tip:** The PFS and PFSMIN options should not be specified together; the PFS option gathers all the information that the PFSMIN option gathers.
PING	Turns on the following trace options: • ARP • ICMP • RAW
POLICY	Trace the stack usage of Policy Rules and Actions.
PORT(list)	Turns on trace by port number.
QUEUE	Turns on trace for stream queue management. **Guideline:** The ENGINE and QUEUE options are aliases. If you turn one on, you turn on all related options, and if you turn one off, you turn off all related options. These alias options are only for recording the trace. When formatting the trace, these options can be filtered separately.
RAW	Turns on trace for the RAW transport protocol.
ROUTE	Trace manipulation of IP Routing Tree.

Table 11. Trace options and groups (continued)

Trace Event	Description
RW	Turns on the following trace options: • ENGINE • PFS • QUEUE • RAW • SOCKET
SERIAL	Turns on trace for lock obtain and release. **Tip:** Using this option slows performance considerably, so use with caution. Also available for the TN3270E Telnet server running in its own address space.
SMTP	Turns on the following trace options: • ENGINE • IOCTL • PASAPI • PFS • QUEUE • SOCKET • TCP
SNMP	Turns on trace for SNMP SET requests.
SOCKAPI	Trace Macro and Call Instruction API calls (see "Socket API traces" on page 74).
SOCKET	Turns on trace for the Sockets API layer.
STORAGE	Turns on trace for storage obtain and release. **Tip:** Using this option slows performance considerably, so use with caution. Also available for the TN3270E Telnet server running in its own address space.
SYSTEM	Turns on the following trace options: • INIT • OPCMDS • OPMSGS • SERIAL • STORAGE • TIMER • WORKUNIT
TC	Turns on the following trace options: • ENGINE • PFS • QUEUE • SOCKET • TCP

Table 11. Trace options and groups (continued)

Trace Event	Description
TCP	Turns on trace for the TCP transport protocol. **Restriction:** The TCP and TCPMIN options should not be specified together; the TCP option gathers all the information that the TCPMIN option gathers. **Slow Performance:** Using this option slows performance considerably, so use with caution.
TCPMIN	Turns on the minimum TCP trace option. **Slow Performance:** The TCP and TCPMIN options should not be specified together; the TCP option gathers all the information that the TCPMIN option gathers. The same is also true for the PFS and PFSMIN options.
TELNET	Turns on trace for TELNET events. Only useful when used by the TN3270E Telnet server.
TELNVTAM (an alias for TELNET)	Turns on trace for TELNET events.
TIMER	Turns on trace for TCP timers. **Slow Performance:** Using this option slows performance considerably, so use with caution. Also available for the TN3270E Telnet server running in its own address space.
TN	Turns on the following trace options for TCP: • PFS • TCP Turns on the following trace option for the TN3270E Telnet server running in its own address space: • TELNET
UD	Turns on the following trace options: • ENGINE • PFS • QUEUE • SOCKET • UDP
UDP	Turns on trace for UDP transport protocol. **Slow Performance:** Using this option slows performance considerably, so use with caution.
VTAM	Shows all of the nondata-path signaling occurring between IF and VTAM.
VTAMDATA	Shows data-path signaling between IF and VTAM, including a snapshot of media headers and some data. **Slow Performance:** Using this option slows performance considerably, so use with caution.
WORKUNIT	Turns on trace for work unit scheduling.
XCF	Turns on trace for XCF events.

Specifying trace options after initialization

After TCP/IP or Telnet initialization, you must use the TRACE CT command to change the component trace options. Each time a new component trace is initiated, all prior trace options are turned OFF, and the new traces are activated.

You can specify the TRACE CT command with or without the PARMLIB member.

You can use IBM Health Checker for z/OS to check whether TCP/IP Event Trace (SYSTCPIP) is active with options other than the default options (MINIMUM, INIT, OPCMDS, or OPMSGS). For more details about IBM Health Checker for z/OS, see Appendix C, "IBM Health Checker for z/OS," on page 999.

Additional filters for SYSTCPIP

The following additional trace filters for limiting the volume of trace data are available:

- The IPADDR keyword filters by IP address
- The PORT keyword filters by port number
- The EID keyword filters by event identifier

  The EID keyword specifies up to 16 trace event identifiers. Each identifer is 8 hexadecimal characters. For example: EID(00010001,00090001,40030003). Use the EID keyword only with the direction of IBM service personnel.

To execute a trace on a particular IP address, use the IP address, port number, ASID, and JOBNAME as targets for filtering the records.

To use this function, start by issuing the TRACE command:

```
TRACE CT,ON,COMP=SYSTCPIP,SUB=(tcpip_procedure_name)
 R 01,OPTIONS=(IPADDR(12AB:0:0:CD30::/60),PORT(1012))
 R 02,OPTIONS=(ENGINE,PFS),END
```

Trace records of type ENGINE or PFS for an IP address of 12AB:0:0:CD30::/60 and a port number of 1012 are captured. The IP address used is the foreign session partner IP address. The port number is the local session partner port number.

When filters are used, the trace record must be accepted by each filter. Each filter can specify multiple values (up to 16), and the trace record must match one of the values.

Table 12 lists the data types and corresponding description.

Table 12. Data types

Data type	Description
Inbound	Data received at the IP layer is considered inbound data. The source IP address and the destination port number are used.
Outbound	Data sent in the PFS layer is considered outbound data. The destination IP address and the source port number are used.

The following are five criteria for selecting trace records for recording:

- TYPE
- JOBNAME
- ASID

- IPADDR
- PORT

Each criterion can specify one or more values. If a criterion has been specified, the record to be traced must match one of the values for that criterion. If a criterion has not been specified, the record is not checked and does not prevent the record from being recorded. However, the record must match all specified criteria.

In the above example, JOBNAME and ASID were not specified, so the value of JOBNAME and ASID in the record are not checked.

Restriction: IPADDR and PORT are exceptions. Some trace records do not have a IP address or a port number. Therefore, the IP address in the trace record is only checked if it is nonzero, and the port number is checked only if it is nonzero.

You can also specify a range of IP addresses to trace. For example,

```
TRACE CT,ON,COMP=SYSTCPIP,SUB=(TCPIP_PROC_NAME)
  R xx, OPTIONS = (IPADDR(nn.nn.nn.nn,{nn.nn.nn.nn/mm.mm.mm.mm}),PORT(pppp{,pppp}))
```

IPADDR
An IP address. Up to 16 addresses can be specified. IPv4 addresses are in dotted decimal notation, for example: 192.48.24.57. IPv6 addresses are in colon-hexadecimal notation or in a combination of both colon-hexadecimal and dotted decimal for IPv4-mapped IPv6 addresses, for example: beef::c030:1839. Use an IP address of 0 for trace records that do not have an IP address. A subnet mask is indicated by a slash (/) followed by the prefix length in decimal or by a dotted decimal subnet mask for IPv4 addresses. The prefix length is the number of one bits in the mask. For IPv4 addresses it might be in the range of 1–32; for IPv6 addresses it might be in the range of 1–128, for example: 192.48.24/24 or 2001:0DB8::0/10, respectively.

PORT
The list of port numbers to be filtered. Up to 16 port numbers can be specified. The port numbers, specified in decimal, must be in the range 0–65535. A trace record with a zero port number is not subject to port number filtering.

You can specify the IPADDR and PORT keywords multiple times in an OPTIONS string. If you do, all the values are saved.

Restriction: All the values in the OPTIONS keyword must be specified in one trace command. The next trace command with an OPTIONS keyword replaces all the options specified.

Formatting event trace records for TCP/IP stacks and Telnet

You can format event trace records using IPCS panels or a combination of IPCS panels and the CTRACE command. For a description of the relevant IPCS panels, see “Steps for formatting component traces using IPCS panels” on page 54.

For more information about other CTRACE options, see the z/OS MVS IPCS Commands.

When using an IPCS panel, enter the trace types in the following format:

```
option DUCB() CID()
```

Following is the syntax for the CTRACE command for TCP/IP stacks and Telnet. For more information about the command and IPCS, see the z/OS MVS IPCS User's Guide.

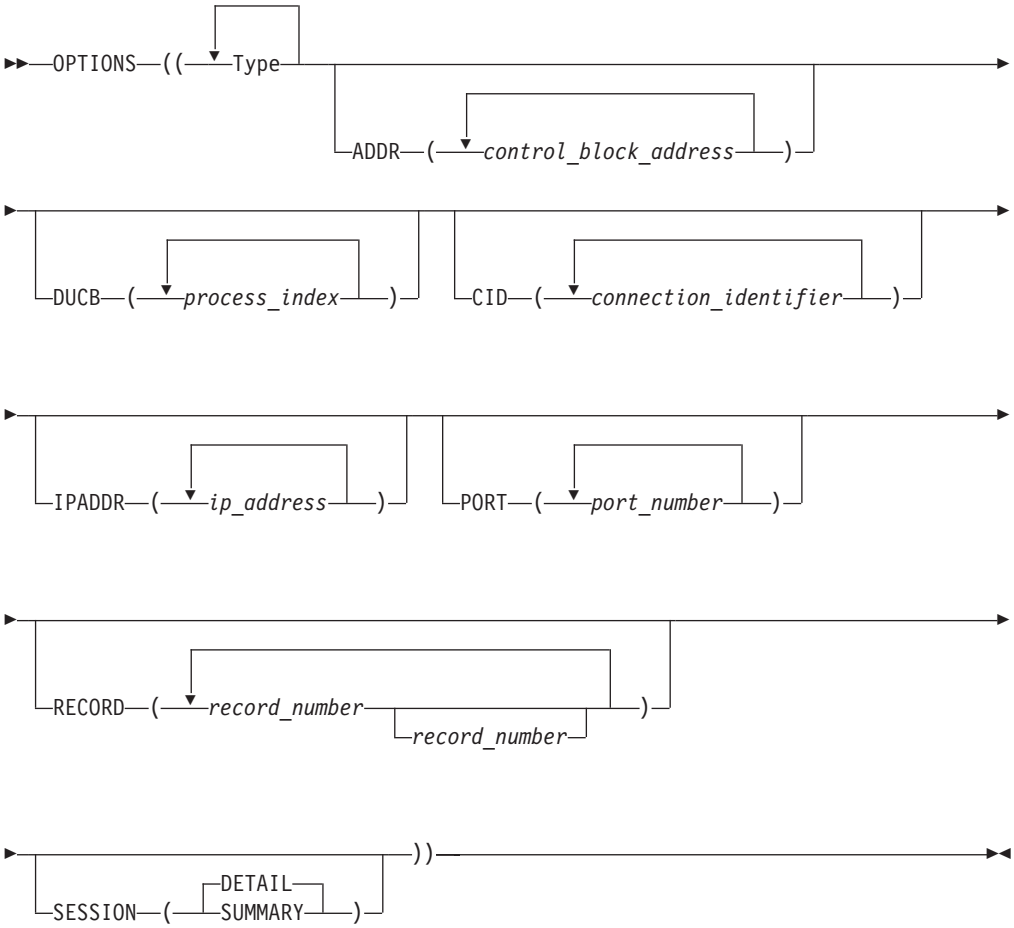

Type Name
The name of a trace type. Only records of these types are formatted. For a list of types, see Table 11 on page 61.

ADDR
A control block address. Up to 16 control block addresses can be specified. Addresses in hexadecimal should be entered as `x'hhhhhhhh'`.

DUCB
A process index for the thread of execution. Up to 16 indexes can be specified. The DUCB index values can be entered either in decimal (such as `DUCB(18)`) or hexadecimal (such as `DUCB(X'12')`), but are displayed in hexadecimal format.

CID
A connection identifier. Up to 16 identifiers can be specified. The CID values can be entered in either decimal (such as `CID(182)`) or hexadecimal (such as `CID(X'0006CE7E')`), but are displayed in hexadecimal. This is the same value that appears in the NETSTAT connections display.

IPADDR
An IP address. Up to 16 addresses can be specified. IPv4 addresses are in

dotted decimal notation, for Example: 192.48.24.57. IPv6 addresses are in colon-hexadecimal notation or in a combination of both colon-hexadecimal and dotted decimal for IPv4-mapped IPv6 addresses, for example: beef::c030:1839. Use an IP address of 0 for trace records that do not have an IP address. A subnet mask is indicated by a slash (/) followed by the prefix length in decimal or by a dotted decimal subnet mask for IPv4 addresses. The prefix length is the number of one bits in the mask. For IPv4 addresses it might be in the range of 1 - 32; for IPv6 addresses it might be in the range of 1–128, for example: 192.48.24/24 or 2001:0DB8::0/10

PORT
A port number. Up to 16 port numbers can be specified. The port numbers can be entered in decimal, such as `PORT(53)`, or hexadecimal, such as `PORT(x'35')`, but are displayed in decimal. These are port numbers in the range 0–65535. Use a port number of 0 for trace records that do not have a port number.

RECORD
The record number can be specified as a single hexadecimal value (for example, `x'hhhhhhhh'`) or as a range (for example, `x'hhhhhhh':x'hhhhhhhh'`). The record number is assigned as the records are written and can be found on the line of equal signs (=) that separates each record.

SESSION(DETAIL|SUMMARY)
Generate a report that shows TCP and socket send/receive activity on a per-session basis. This formatting option operates on TCPMIN and PFSMIN trace entries.

Tip: The SESSION formatting option for SYSTCPIP traces is useful for analyzing TCP traffic that is flowing over an IPSec tunnel. If you are accustomed to using packet trace (SYSTCPDA) for diagnosis, IPSec encryption of TCP headers and insertion of AH|ESP headers can result in an inability to do TCP-layer diagnosis (for example, matching up ACKs with data packets, spotting retransmissions, analyzing flow control issues) with the packet trace report. By contrast, the SESSION formatting option for SYSTCPIP trace operates on TCP headers in the clear before encryption and after decryption, so this report can be more informative than a standard packet trace if encryption is involved or AH|ESP headers are present.

DETAIL
List each inbound or outbound flow within the socket and TCP layers. If neither the PFSMIN or TCPMIN CTRACE options were enabled in the trace, a null report is produced. DETAIL is the default.

SUMMARY
Show only the summary statistics.

Standard TSO syntax is used for the keywords and their values. For example, CID (1 2 3).

Figure 13 on page 71 shows the beginning of the CTRACE formatted output. The CTRACE command parameters are followed by the trace date and column headings. Then, there is one TCP/IP CTRACE record with four data areas.

```
COMPONENT TRACE FULL FORMAT
     COMP(SYSTCPIP)SUBNAME((TCPSVT))

     **** 11/03/1999
     SYSNAME   MNEMONIC  ENTRY ID    TIME STAMP     DESCRIPTION
     -------   --------  --------  --------------- -------------
1     VIC142    PFS       60010018  14:57:59.207826   Socket IOCTL Exit
2    HASID..001E      PASID..000E      SASID..001E      USER...OMPROUTE
3    TCB....007E7A68  MODID..EZBPFIOC  REG14..161D86C0  DUCB...0000000C
4    CID....0000003A  PORT.....0
   IPADDR. 3F98::D002:A521
5      ADDR...00000000  14D9EED0  LEN....000000A0  OSI
6         +0000  D6E2C940  000000A0  00000000  00000000  | OSI ............ |
          +0010  0500001B  14D9EF70  00500AC8  00000000  | .....R...&.H.... |
          +0020  00000000  00000000  00000000  00000000  | ................ |
          +0030  00000000  00000000  00000000  00281080  | ................ |
          +0040  14D9FC0C  00000C00  14D9FFE8  00000000  | .R.......R.Y.... |
          +0050  00000000  00000000  00000000  00000000  | ................ |
          +0060  00000000  00000000  00000000  00000000  | ................ |
          +0070  00000000  00000000  00000000  00000000  | ................ |
          +0080  00000000  00000000  00000000  00000000  | ................ |
          +0090  00000000  00000000  00000000  00000000  | ................ |
       ADDR...00000000  12D7F874  LEN....00000004  SCB Flags
          +0000  00280000                                | ....             |
       ADDR...00000000  12E88598  LEN....00000010  Return Value Errno ErrnoJr
          +0000  C5D9D9D5  FFFFFFFF  00000462  740E006B  | ERRN............, |
       ERRNO..-1,     462,  740E006B
       ADDR...00000000  14D9F4E4  LEN....00000048  IOCTL Request
          +0000  C3C6C7D4  D9C5D840  0000008E  00000462  | CFGMREQ ........ |
          +0010  00000320  00000500  00000000  00000000  | ................ |
          +0020  740E0005  00000000  14B4C7C0  00000000  | .........G{.... |
          +0030  00000000  00050063  00000000  00000000  | ................ |
          +0040  F3F1F0F1  00000000                      | 3101....         |
7    ================================================================0000573E
```

Figure 13. Start of component trace full format

The parts of the TCP/IP CTRACE record are:

1 Standard IPCS header line, which includes the system name (VIC142), TCP/IP option name (PFS), time stamp, and record description.

2 TCP/IP header line with address space and user (or job name) information.

3 TCP/IP header line with task and module information.

4 TCP/IP header line with session information (CID, IP address, and port number).

5 TCP/IP header line for a data area. This line has the address (first four bytes are the ALET), the length of data traced, and the data description. Following the description, the actual data is in dump format (hexadecimal offset, hexadecimal data, and EBCDIC data).

6 There are four data areas in this example. The third data area (Return Value Errno ErrnoJr" has an extra line. The ERRNO line is added only when the return value is -1 and the ERRNO indicates an error. In this example, the return code is hexadecimal 462 (decimal 1122). See z/OS Communications Server: IP and SNA Codes for more information.

7 TCP/IP trailer and separator line with the record sequence number (hexadecimal 573E).

Additional fields in CTRACE output

The ERRNO line in Figure 13 is one of two cases in which the formatter extracts data and formats it in a special way. The other case is for "TCB CTRL" and "IUDR" data. Several fields are copied from the data and formatted with character

interpretation of fields, such as converting values to decimal or dotted decimal. Figure 14 is an example. Note the additional fields (TcpState, TpiState, and others) following the hexadecimal data.

```
BOTSWANA  TCP       40030002  20:51:35.652462   Select/Poll Exit Detail
HASID..0082      PASID..0088      SASID..000E      USER...POLAGENT
TCB....007E4640  MODID..EZBTCFSP  REG14..10FD7C5E  DUCB...00000016
CID....000004DC  PORT....1925
IPADDR. 197.011.106.001
  ADDR...00000000  116B04DC  LEN....00000004  Select function code
     +0000  00000002                                | ....            |
  ADDR...00000000  116B0668  LEN....00000004  Output condition indicators
     +0000  40000000                                |  ...            |
  ADDR...00000000  7F60C508  LEN....000003D8  Transmission Control Block
     +0000  E3C3C240  C3E3D9D3  00050009  81801000  | TCB CTRL....a... |
     +0010  00000000  00000000  00000000  138C4F08  | ..............|. |
      ...
     +0170  00000000  00020000  00003000  45000028  | ................ |
     +0180  1CB14000  40060000  C50B6A01  C50B6A01  | .. . ...E...E... |
     +0190  00000000  00000000  00000000  00000000  | ................ |
     +01A0  00000000  00000000  00000000  00000000  | ................ |
     +01B0  00000000  00000000  0000FFFF  FFFF4000  | .............. . |
     +01C0  00000000  00000000  00000000  00000001  | ................ |
     +01D0  07850185  F4258CA0  F425A310  50107F32  | .e.e4...4.t.&.". |
     +01E0  00000000  0004FFCB  01030300  0101080A  | ................ |
      ...
     +03D0  010E1301  0E21010E                      | ........         |
  TcpState..ESTAB     TpiState..WLOXFER
  SrcPort..1925       SrcIPAddr. 197.11.106.1
  DstPort..389        DstIPAddr. 197.11.106.1
  FLAGS.....ACK
```

Figure 14. Component trace full format showing character interpretation of fields

SESSION report example

Purpose

This report shows traffic for TCP sessions, and is generated from PFSMIN and TCPMIN event trace entries.

Format

```
CTRACE COMP(SYSTCPIP) SUB((TCPCS)) SHORT OPTIONS((SESSION))
```

Example

```
COMPONENT TRACE SHORT FORMAT
 COMP(SYSTCPIP)SUBNAME((TCPCS))
 OPTIONS((SESSION))

=================================================================================================================
Performance Sessions Report
     16 Sessions Found

Explanation of Unique Fields
=================================================================================================================
RECTYPE  | UNIQUE-1 | UNIQUE-2 | UNIQUE-3 | UNIQUE-4 | UNIQUE-5 | UNIQUE-6 | UNIQUE-7 | UNIQUE-8
---------+----------+----------+----------+----------+----------+----------+----------+------------+
PFSE     |          |          |          |          |          |SCB Flags |OSI Flags |Sel Bits
PFSX     |RetVal    |RetCode   |RsnCode   |          |          |SCB Flags |OSI Flags |Sel Bits
TCRE     |          |Msg Flags |Event Ptr |UsrRecvBuf|Recv QLen |Recv Win  |OSI Flags |Conn State
TCRX     |RetVal    |RetCode   |RsnCode   |UsrRecvBuf|Recv QLen |RecvWin   |          |Conn State
TCWE     |          |Msg Flags |Event Ptr |UsrSendBuf|Send QLen |Send Win  |OSI Flags |Conn State
TCWX     |RetVal    |RetCode   |RsnCode   |UsrSendBuf|Send QLen |Send Win  |          |Conn State
TCRV     |IpID/TCPFl|Seq Num   |Ack Num   |Bytes Rcvd|Recv QLen |Recv Win  |          |
TCSN     |IpID/TCPFl|DataInPipe|NumPktSent|Bytes Sent|Send QLen |Send Win  |Congst Win|Xmit Cnt/Tot
TCSQ     |          |          |          |          |          |          |          |
TCRD     |DropRsn   |          |          |          |          |          |          |
---------+----------+----------+----------+----------+----------+----------+----------+------------+
```

```
*Note:  If the trace does not contain TCPMin entries, the Session Summary will be missing or incorrect.
Directional Arrow Flow Descriptions:
-->  Application to Network Call
<--  Network to Application Call
S->  Application Send/Write Call
<-S  Applicaion Send/Write Return
R->  Application Receive
<-R  Application Receive Return
--Q  Application Send/Write Data Queued (slow start, congestion avoidance, delayack, etc)
--D  Inbound Data Dropped (malformed, bad checksum, no listener, etc)

================================================================================================================================
TCPIP SESSION FLOW FOR CID:    0000002A
  Number of trace entries:     206394
--------------------------------------------------------------------------------------------------------------------------------
Application Information:               Local               Remote
  IP address:                          9.67.170.70         9.67.170.71
  Port:                                2121                1025
  Jobname:                             AWM4
  Asid:                                0034

Connection Information:
  Total Bytes Transferred:             248717618
    Bytes Sent:                        13
    Bytes Received:                    248717605
  Total number of retransmissions:     0
    Retransmissions sent:              0
    Retransmissions received           0
  Queue lengths
    Largest send queue length:         1
    Largest receive queue length:      115179
  Window sizes
    Congestion window information
      Largest congestion window:       91806
      Smallest congestion window:      73723
     Receive window information
       Largest receive window:          367261
       Smallest receive window:         225745
     Send window information
       Largest send window:             368640
       Smallest send window:            368632
 Last known connection state is ESTABLISHED

 Time Frame Information:
   First Timestamp:                     16:55:23.079072
   Last Timestamp:                      16:56:13.697196
   Duration of session:                 00:00:50.618123
   Maximum Delta:                       00:00:01.133860
      (at Timestamp):                   16:56:10.105556

 Session Summary:
 RECTYP| TIME            |PFS|TCP|IP | MODID   UNIQUE-1 UNIQUE-2 UNIQUE-3 UNIQUE-4 UNIQUE-5 UNIQUE-6 UNIQUE-7  UNIQUE-8
 ------+-----------------+---+---+---+------  --------- -------- -------- -------- -------- -------- -------- ---------
 TCRX  | 16:55:23.079072 |   |<--|   | TCFRD   0000FFFF 00000000 00000000 0000FFFF 00000000 0004A001              0005
 PFSX  | 16:55:23.079074 |<-R|   |   | PFSDR   0000FFFF 00000000 00000000                   90280000     0000      0000
 PFSE  | 16:55:23.079082 |R->|   |   | PFSDR                                                90280000     0000      0000
 TCRE  | 16:55:23.079083 |   |-->|   | TCFRD            00000000 00000000 0000FFFF 00000000 0004A001     0000      0005
 TCRV  | 16:55:23.081066 |   |   |<--| TCRD   C7A3/0010 C4EC5B70 8AD4F838 00000563 00000000 0004A001
 TCRV  | 16:55:23.081069 |   |   |<--| TCRD   C7A4/0010 C4EC60D3 8AD4F838 00000563 00000000 0004A001
 TCRV  | 16:55:23.081070 |   |   |<--| TCRD   C7A5/0010 C4EC6636 8AD4F838 00000563 00000000 0004A001
 TCRV  | 16:55:23.081071 |   |   |<--| TCRD   C7A6/0010 C4EC6B99 8AD4F838 00000563 00000000 0004A001
 TCRV  | 16:55:23.081073 |   |   |<--| TCRD   C7A7/0010 C4EC70FC 8AD4F838 00000563 00000000 0004A001
 TCRV  | 16:55:23.081074 |   |   |<--| TCRD   C7A8/0018 C4EC765F 8AD4F838 000002D2 00000000 0004A001
 TCSN  | 16:55:23.081078 |   |   |-->| TCSND  E49D/0010 00000000 00000001 00000000 00000000 00059FF8 00011FFB 0000/0000
 TCRV  | 16:55:23.090273 |   |   |<--| TCRD   C779/0010 C4EB7932 8AD4F838 00000563 00000000 0004A001
 TCRV  | 16:55:23.090281 |   |   |<--| TCRD   C77A/0010 C4EB7E95 8AD4F838 00000563 00000563 00059A9D
 TCRV  | 16:55:23.090283 |   |   |<--| TCRD   C77B/0010 C4EB83F8 8AD4F838 00000563 00000AC6 0005953A
 TCRV  | 16:55:23.090291 |   |   |<--| TCRD   C77C/0010 C4EB895B 8AD4F838 00000563 00001029 00058FD7
 TCRV  | 16:55:23.090292 |   |   |<--| TCRD   C77D/0010 C4EB8EBE 8AD4F838 00000563 0000158C 00058A74
 TCRV  | 16:55:23.090294 |   |   |<--| TCRD   C77E/0010 C4EB9421 8AD4F838 00000563 00001AEF 00058511
 TCRV  | 16:55:23.090295 |   |   |<--| TCRD   C77F/0010 C4EB9984 8AD4F838 00000563 00002052 00057FAE
 TCRV  | 16:55:23.090297 |   |   |<--| TCRD   C780/0010 C4EB9EE7 8AD4F838 00000563 000025B5 00057A4B
 TCRV  | 16:55:23.090298 |   |   |<--| TCRD   C781/0010 C4EBA44A 8AD4F838 00000563 00002B18 000574E8
 TCRV  | 16:55:23.090300 |   |   |<--| TCRD   C782/0010 C4EBA9AD 8AD4F838 00000563 0000307B 00056F85
 TCRV  | 16:55:23.090301 |   |   |<--| TCRD   C783/0010 C4EBAF10 8AD4F838 00000563 000035DE 00056A22
 TCRV  | 16:55:23.090303 |   |   |<--| TCRD   C784/0010 C4EBB473 8AD4F838 00000563 00003B41 000564BF
 TCRV  | 16:55:23.090304 |   |   |<--| TCRD   C785/0010 C4EBB9D6 8AD4F838 00000563 000040A4 00055F5C
 TCRV  | 16:55:23.090306 |   |   |<--| TCRD   C786/0010 C4EBBF39 8AD4F838 00000563 00004607 000559F9
 TCRV  | 16:55:23.090307 |   |   |<--| TCRD   C787/0010 C4EBC49C 8AD4F838 00000563 00004B6A 00055496
```

```
TCRV  | 16:55:23.090309 |    |    |<--| TCRD   C788/0010 C4EBC9FF 8AD4F838 00000563 000050CD 00054F33
TCRV  | 16:55:23.090310 |    |    |<--| TCRD   C789/0010 C4EBCF62 8AD4F838 00000563 00005630 000549D0
TCRV  | 16:55:23.090312 |    |    |<--| TCRD   C78A/0010 C4EBD4C5 8AD4F838 00000563 00005B93 0005446D
```

Socket API traces

The SOCKAPI option, for the TCP/IP CTRACE component SYSTCPIP, is intended to be used for application programmers to debug problems in their applications. The SOCKAPI option captures trace information related to the socket API calls that an application might issue. The SOCKET option is primarily intended for use by TCP/IP Service and provides information meant to be used to debug problems in the TCP/IP socket layer, UNIX System Services, or the TCP/IP stack.

CTRACE is available only to users with console operator access. If the application programmer does not have console access, someone must provide the CTRACE data to the programmer. For security reasons, it is suggested that only the trace data related to the particular application be provided. The following sections explain how to obtain the trace data for a particular application, format it, and save the formatted output. The application data can be isolated when recording the trace, or when formatting it, or both.

z/OS provides several socket APIs that applications can use. The figure in the example shows different APIs along with the high level flows of how they interact with the TCP/IP stack.

The SOCKAPI trace output is captured in the Sockets Extended Assembler Macro API (the Macro API). Given the structure of the TCP/IP APIs, this trace also covers the Call Instruction API, the CICS Socket API, and the IMS socket API. Some of the socket APIs based on the Macro API currently encapsulate some of the Macro API processing.

For example, in a CICS TS environment, CICS sockets-enabled transactions do not have to issue an SOCKAPI call. Rather, this is done automatically for the socket API by the TCP/IP CICS TRUE (Task Related User Exit) component layer. If the socket API trace is active, trace records for the SOCKAPI calls are created.

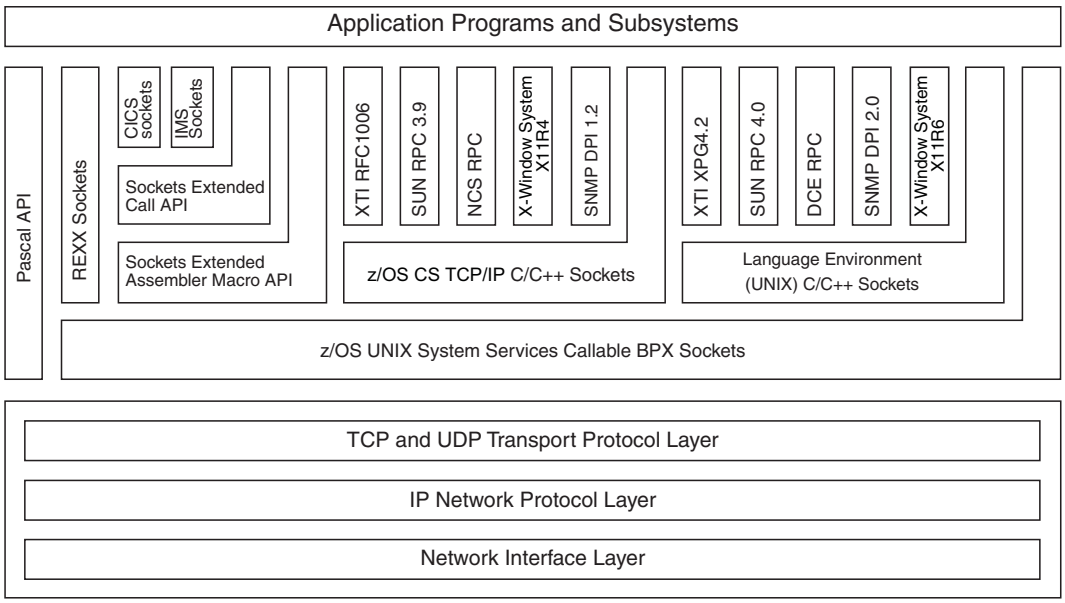

Figure 15. TCP/IP networking API relationship on z/OS

Recommended options for the application trace

The CTRACE facility has flexibility such as filtering, combining multiple concurrent applications and traces, and using an external writer.

Guidelines: Consider the following guidelines when using CTRACE:

- Although the CTRACE can be used to trace multiple applications at the same time and in conjunction with other trace options, it is not recommended. Multiple traces make problem determination more difficult.
- For performance reasons, the data being recorded should be filtered, to minimize the overhead of recording the trace, to make formatting faster, to save storage, and to minimize wrapping (overwriting of older trace records by new trace records).

Ideally, you should use the CTRACE facility to capture all the SOCKAPI trace records for one application. The trace can be filtered various ways when formatting. If necessary, you can limit the trace data collected by IP address or port number, but you risk some records not being captured. For example, the problem might be that the wrong IP address or port number was coded or used. Both the IP address and port number are formatting options.

Guidelines: Consider the following guidelines to optimally capture the application data:

- Trace only one application.

  Use the job name or ASID option when capturing the trace to limit the trace data to one application.
- Trace only the SOCKAPI option.

  To get the maximum number of SOCKAPI trace records, specify only the SOCKAPI option.

  Tip: You also receive exception records. Exception records are always traced because they are considered unusual events.

- Use an external writer.

  The external writer is recommended to:
  - Separate the SOCKAPI trace records from other internal data that exist in a dump (for security and other reasons)
  - Avoid interrupting processing with a dump of the trace data
  - Keep the buffer size from limiting the amount of trace data
  - Avoid increasing the buffer size, which requires restarting TCP/IP
  - Handle many trace records
- Trace only one TCP/IP stack.

  If you are running with multiple TCP/IP stacks on a single z/OS image, use the external writer for only one TCP/IP stack.
- Activate the data trace only if more data is required.

  The SOCKAPI trace contains the first 96 bytes of data sent or received, which is sufficient. If additional data is needed, the data trace records can be correlated with the SOCKAPI records.

Collecting the SOCKAPI trace option

This topic describes how to collect the trace for use by application programmers.

The existing CTRACE facility for TCP/IP's SYSTCPIP component is used for the SOCKAPI trace option. Collecting the trace is described generally in "Component trace" on page 45.

The trace can be started automatically when TCP/IP starts or can be started or modified while TCP/IP is executing. A CTRACE PARMLIB member is required for starting the trace automatically, and can optionally be used after TCP/IP has been started.

CTRACE PARMLIB member CTIEZBxx:
Sample member CTIEZB00 is shipped with TCP/IP.

TCP/IP start procedure:
The CTRACE PARMLIB member can be specified in the TCP/IP start procedure or on the START command. The sample TCPIPROC start procedure specifies member name CTIEZB00. Specifying the member name on the START command depends on how the TCP/IP start procedure is coded.

The following example illustrates overriding the PARMLIB member name using the sample TCPIPROC start procedure.

```
S TCPIPROC,PARM='CTRACE(CTIEZBAN)'
```

Use the `TRC` option to specify the suffix of the SYS1.PARMLIB member for SYSTCPIP CTRACE initialization. The `TRC` option appends the two letters to CTIEZB. The full member name is CTIEZBxx. The default value is 00. In this example, the PARMLIB member for SYSTCPIP is CTIEZBAN, an equivalent command is

```
S TCPIPROC,PARM='TRC=AN'
```

Use the `IDS` option to specify the suffix of the SYS1.PARMLIB member for SYSTCPIS CTRACE initialization. The `IDS` option appends the two letters to CTIIDS. The full member name is CTIIDSxx. The default value is 00.

```
S TCPIPROC,PARM='IDS=AN'
```

You can specify multiple parameters. If you specify both the CTRACE and TRC parameters, the parameter that appears last in the parameter string is used.

TRACE command:
Use the MVS TRACE command to start, modify, or stop the trace after TCP/IP has been started. The TRACE command replaces all prior settings except the buffer size. When modifying the options, be sure to specify the SOCKAPI option.

The examples below show how to start the trace.

The SUB option is the subtrace name, which for TCP/IP, is the job name of the stack (this is the TCP/IP start procedure name). In the following examples, the subtrace is TCPIPROC (the name of the sample procedure), and the variable fields are in lowercase.

To activate the trace with just the SOCKAPI option, code the following:

```
TRACE CT,ON,COMP=SYSTCPIP,SUB=(tcpiproc)
R n,JOBNAME=(ezasokjs),OPTIONS=(sockapi),end
```

To specify a PARMLIB member, which contains the trace options, code the following:

```
TRACE CT,ON,COMP=SYSTCPIP,SUB=(tcpiproc),PARM=ctiezban
```

To stop the trace, either use the TRACE CT,OFF command or reissue the TRACE CT,ON command with different parameters.

The following is an example of the OFF option:

```
TRACE CT,OFF,COMP=SYSTCPIP,SUB=(tcpiproc)
```

When using the TRACE command, be sure to notice message ITT038I, which indicates whether the command was successful or not. The following is an example of ITT038I:

```
  14.11.29  ITT038I NONE OF THE TRANSACTIONS REQUESTED VIA THE TRACE CT
  COMMAND WERE SUCCESSFULLY EXECUTED.
or
  14.11.40  ITT038I ALL OF THE TRANSACTIONS REQUESTED VIA THE TRACE CT
  COMMAND WERE SUCCESSFULLY EXECUTED.
```

See z/OS MVS System Commands for more information about the TRACE command.

External writer:
If the trace is active, it is always written to an internal buffer (whose size is set to BUFSIZE during TCP/IP initialization). The internal buffer is available only in a dump of TCP/IP. Optionally, the trace can also be written to an external data set using the MVS CTRACE external writer. If you use an external writer, the trace records are copied to a data set.

To use an external writer, you must create a procedure that specifies the job to run (the external writer) and the trace output data sets. Also, see z/OS MVS Diagnosis: Tools and Service Aids for more information about CTRACE, the external writer (including a sample procedure), dispatching priority for the external writer job, and wrapping.

The external writer must be started before the trace can be activated. The trace must be inactivated before the writer can be stopped. The writer must be stopped

before the data set can be formatted or transferred. For example, here is a sequence of commands for using an external writer procedure named ctw:

```
TRACE CT,WTRSTART=ctw
 TRACE CT,ON,COMP=SYSTCPIP,SUB=(tcpiproc)
   R n,JOBNAME=(ezasokjs),OPTIONS=(sockapi),WTR=ctw,end

 <run application being traced>

 TRACE CT,OFF,COMP=SYSTCPIP,SUB=(tcpiproc)
 TRACE CT,WTRSTOP=ctw
```

The external data set (specified in the procedure "ctw") is now available for formatting.

Filtering options when recording the trace:
Options for filtering include the following:

Component
Required - SYSTCPIP for SOCKAPI.

Subtrace
Required - TCP/IP stack name.

Trace option
Highly recommended to limit the tracing to the SOCKAPI option. You can also filter on this option when formatting the trace.

Jobname
Highly recommended for socket applications to limit the trace to one application. You can also filter on this option when formatting the trace.

ASID Highly recommended as an alternative to the job name if the application has already started running (otherwise, the ASID is unknown). You can also filter on this option when formatting the trace.

IP address
Recommended only for certain scenarios (see discussion below). The IP address is a filtering option when formatting the trace.

Port Recommended only for certain scenarios (see discussion below). The port number is a filtering option when formatting the trace.

If trace data for multiple applications is collected in the same data set or in a dump, the trace output should be filtered so that application programmers see only the data for their applications for security reasons.

Use the IP address and Port options to filter the trace, both when collecting the trace and when formatting the trace. Generally, it is best to collect all the application records to avoid having to re-create the problem. After the records are collected, you can filter the records various ways when formatting the trace.

An example scenario in which you would only want to collect records for one IP address is if there is a problem with a particular remote client, and the local application has many clients. If you tried to record the trace records for all clients, there could be a lot of data and the trace could wrap, thus overwriting older records. Note that if you specify an IP address when collecting the trace, the trace records with no IP address are also collected. So you get all the records for the problem client, and some other client records.

An example scenario, in which you would only want to collect records for one port number, is if there is a problem with a server on one port. If you specify a port

number when collecting the trace, the trace records with no port number are also collected. You get all the records for the problem server application, and some other applications' records.

IP address/port filtering, when specified, has a varying effect depending on the type of socket call being traced. Table 13 describes the effect of IP address/port filtering for the different types of socket API calls. The Yes or No specified in columns 2 and 3 indicates whether local port filtering and remote IP address filtering can be activated for the socket calls in column 1. Yes means that if a filter is set, only the calls matching that filter are collected. No means that whether a filter is specified, all the calls are collected (no filtering is done).

Table 13. IP address and port filtering effect on different types of socket API calls

Socket call	**Filtering active?**	
	Local port	**Remote IP address**
ACCEPT	Yes	No (1)
BIND	Yes/No (2)	No
CONNECT	Yes/No (3)	Yes
CANCEL FREEADDRINFO GETADDRINFO GETCLIENTID GETHOSTBYADDR GETHOSTBYNAME GETHOSTID GETHOSTNAME GETNAMEINFO INITAPI RECVFROM RECVMSG SELECT SELECTEX SENDMSG SENDTO SOCKET TAKESOCKET TERMAPI	No	No
LISTEN	Yes	No

Table 13. IP address and port filtering effect on different types of socket API calls (continued)

Socket call	Filtering active?	
	Local port	Remote IP address
CLOSE	Yes	Yes
GETPEERNAME		
GETSOCKNAME		
GETSOCKOPT		
GIVESOCKET		
FNCTL		
IOCTL		
READ		
READV		
RECV		
SHUTDOWN		
SEND		
SETSOCKOPT		
WRITE		
WRITEV		

Where Yes is indicated in Table 13 on page 79, the assumption is made that the information necessary for the filtering option is available. For example, if a SEND is issued on a socket that is not bound or not connected, no filtering takes place. In addition, the following describe some of the special considerations for the different socket calls in the previous table.

1. Even though the remote IP address is available after an ACCEPT call, it is not used for filtering the exit ACCEPT trace record. This is done to avoid confusion where the entry trace record for ACCEPT would not be filtered, but the exit trace record would.
2. Assumes a BIND issued for a nonzero port. If a BIND is issued for port 0 (meaning an ephemeral port is assigned by TCP/IP), no filtering takes place for this BIND call.
3. If the socket is bound at the time of the CONNECT, local port filtering is honored. Otherwise, the CONNECT is not subject to local port filtering.

Monitoring the trace:
Use the MVS DISPLAY TRACE command to check the trace options currently in effect. The following example is an example of a console showing the display command and the resulting output (the line numbers were added for discussion reference).

```
 1.     14.27.14  D TRACE,COMP=SYSTCPIP,SUB=(tcpiproc)
 2.     14.27.14  IEE843I 14.27.14  TRACE DISPLAY
 3.             SYSTEM STATUS INFORMATION
 4.      ST=(ON,0064K,00064K) AS=ON  BR=OFF EX=ON  MT=(ON,064K)
 5.       TRACENAME
 6.       =========
 7.       SYSTCPIP
 8.                            MODE BUFFER HEAD SUBS
 9.                            ====================
10.                            OFF         HEAD    1
11.          NO HEAD OPTIONS
12.       SUBTRACE              MODE BUFFER HEAD SUBS
13.      ------------------------------------------------------
```

```
14.       TCPIPROC            ON   0008M
15.          ASIDS      *NONE*
16.          JOBNAMES   EZASOKJS
17.          OPTIONS    SOCKAPI
18.          WRITER     CTW
```

For component SYSTCPIP, do not be misled by line 10 in the example. It always says the trace is off because TCP/IP uses the subtrace for all tracing. The subtrace TCPIPROC on line 14 indicates the actual state of the trace. In this example, the trace is active (ON) with an internal buffer size of eight megabytes and only the SOCKAPI option is active. Only one application (EZASOKJS) is being traced and the trace is being written to an external writer.

Line **Description**

1 The MVS DISPLAY TRACE command. For more information about this command, see z/OS MVS System Commands.

2–4 These are explained in z/OS MVS System Messages, Vol 1 (ABA-AOM) for IEE843I.

5–7 Show that this is the CTRACE component SYSTCPIP.

8–11 These are not applicable for TCP/IP because TCP/IP uses only the subtrace facility of the MVS CTRACE service. Instead of activating a global trace, the trace options are specified for each stack individually. Thus, there can be multiple TCP/IP stacks with different CTRACE options. Note however that line 10 is useful — it shows that there is one subtrace (meaning one TCP/IP stack is active).

14 Shows the "subtrace" name is the TCP/IP procedure name (TCPIPROC in this example), whether the trace is active (MODE=ON), and the buffer size is eight megabytes. The buffer size is the number of bytes that is used for recording the trace.

15–16 Show the ASID and JOBNAME filtering values. If any ASIDs or JOBNAMEs are listed, only those trace entries matching the ASID or JOBNAME are collected. "ASIDS *NONE*" indicates that all address spaces are being traced (there is no filtering).

17 Shows the specific options that are active, as specified in the TRACE command or in the CTIEZBxx PARMLIB member. If port or IP address filtering were active, they would appear on this line.

18 Shows the external writer is inactive. If the writer is active, the writer procedure name is shown instead of *NONE*.

Capturing the trace:
If you use only the internal buffer, you must obtain a dump of TCP/IP to view the CTRACE records. It is usually a good idea to also capture the application address space. For example, using the MVS DUMP command, type the following commands.

Tip: The SDATA options specified are appended to other options.

The SDATA options shown here are the generally recommended options.

```
DUMP COMM=(Sample dump for SOCKAPI)
R n,JOBNAME=(tcpiproc,ezasokjs),CONT
R n,SDATA=(ALLNUC,CSA,LPA,LSQA,RGN,SWA,SQA,TRT),CONT
R n,END
```

Notes:

- You can type the first three commands in advance, and you can then just type the fourth command at the correct moment to capture the events.
- If you use the external writer, "External writer" on page 77, explains how to capture the trace in a data set.

Formatting the SOCKAPI trace option

Use the IPCS CTRACE command to format the trace, both for a dump and for an external writer. Interactively, you can either type the CTRACE command on the IPCS Command panel or you can use the panel interface. IPCS is also available in batch. Whichever interface you choose, for TCP/IP we recommend using the CTRACE QUERY command to find out what subtraces are contained in the data set. For example, the CTRACE QUERY(SYSTCPIP) SHORT command produced the following output:

```
COMPONENT TRACE QUERY SUMMARY

       COMPONENT SUB NAME
       --------- --------
 0001. SYSTCPIP  TCPSVT
 0002. SYSTCPIP  TCPSVT3
 0003. SYSTCPIP  TCPSVT1
 0004. SYSTCPIP  TCPSVT2
```

There are several filters available that can help to limit the amount of data formatted. In addition to the CTRACE options (start and stop time, and such) provided by IPCS, there are some options specifically for TCP/IP:

DUCB Not applicable for SOCKAPI. (DUCB is an internal TCP/IP token.)

CID (connection identifier)
Not applicable for SOCKAPI.

IPADDR
Use for SOCKAPI. Specify the IPv4 addresses in dotted decimal format, with an optional prefix value (1 to 32) or a subnet mask in dotted decimal form. Specify the IPv6 address in colon-hexadecimal notation (or in a combination of colon-hexadecimal and dotted decimal for IPv4–mapped IPv6 addresses), with an optional prefix value (1 to 128). Several socket calls do not use an IP address. To see the trace records without an IP address (or with an IP address of all zeros), specify zero for one of the IPADDR values. For example, IPADDR(0,9.67.113/24) formats all CTRACE records with an IP address of 000.000.000.000 and formats all CTRACE records with an IP address of 009.067.113.*, where * is any number from 0 to 255.

PORT Use for SOCKAPI. Specify the port number in decimal. Several socket calls do not have an associated port number, such as INITAPI and SOCKET. To see the trace records without a port (or with a port of 0), specify zero for one of the port values. For example, `PORT(0,389,1925)`.

You can save the formatted output to the IPCSPRNT data set.

If the formatted output does not contain the records you expect:

- In a dump, you can check the options specified when recording the trace by using the TCPIPCS TRACE command to display the TCP/IP CTRACE filtering options in effect. This also indicates whether any records were lost. See Chapter 6, "IPCS subcommands for TCP/IP," on page 173 for more information about the TCPIPCS TRACE command.

- For either a dump or an external writer data set, use the CTRACE QUERY command to see what tracing was in effect (subtrace name, start and stop times). For a dump, this command also shows the buffer size and options. For example, the command CTRACE QUERY(SYSTCPIP) SUB((TCPIPROC)) FULL command produced the following output for a dump:

```
COMPONENT TRACE QUERY SUMMARY

 COMP(SYSTCPIP)SUBNAME((TCPIPROC))

     START = 01/10/2000 19:49:21.234490   GMT
     STOP  = 01/10/2000 19:51:51.360653
 Buffer size: 0256K

 OPTIONS: ACCESS  ,OPCMDS  ,OPMSGS  ,QUEUE   ,ROUTE   ,INIT    ,SOCKAPI ,SOCKET

 OPTIONS: MINIMUM
```

  For TCP/IP, the first line of "options" (showing ACCESS) is the applicable one. This shows the options as specified on the command line or in the CTIEZBxx PARMLIB member.

See z/OS MVS IPCS User's Guide for more information about CTRACE formatting. See z/OS MVS IPCS Commands for more information about the CTRACE command.

Reading and interpreting the SOCKAPI trace option

The SOCKAPI trace records trace the input and output parameters for most of the API calls. The API calls not traced are GETIBMOPT, TASK, GLOBAL, NTOP, PTON, and any API calls that fail before the trace point is reached. (An API call fails if module EZBSOH03 cannot be located, if EZBSOH03 is unable to obtain storage, and so on.) In addition to tracing API calls, trace records are created for a few special situations (Default INITAPI and Unsolicited Event exit being driven). For API calls, there is an Entry record describing the input parameters, and an Exit record describing the output parameters (with some input parameters repeated for clarification). For asynchronous calls, there is also an Async Complete (Asynchronous Complete) record (see “Examples of SOCKAPI trace records” on page 85).

The following examples include:

- A SOCKAPI trace record
- Trace records for asynchronous applications
- Resolver API calls
- External IOCTL commands
- API Call with an IOV parameter
- Default INITAPI
- Default TERMAPI
- SELECT
- SELECTEX
- Token error
- Unsolicited event exit

A SOCKAPI trace record:

A typical SOCKAPI record is shown below. This example is a READ Entry.

The lines are numbered for discussion reference only. The description for each line is for the example shown. Lines 1-5 are the separator and header lines that exist for all SOCKAPI trace records. Lines 6-7 are optional header lines.

The parameters for the specific call follow the header lines. For Entry records, the input parameters are shown. For Exit and Asynchronous Complete records, the output parameters are shown and some input parameters might also be shown for reference. Parameters are only formatted if they were specified in the call (optional parameters not supplied are not formatted). The parameters are listed in a specific order for consistency. The parameter names are the same as the names in z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference with a few exceptions; for example, S is formatted as SOCKET. The parameter name, value, and address are shown on one line if the value fits. Numeric parameter values are in decimal unless followed by a lowercase x indicating hexadecimal. Whenever possible, the values are interpreted (such as ERRNO) for reference.

```
 1.  ===================================================================00007FE8
 2.   MVS026    SOCKAPI   60050042  19:31:08.338135   READ Entry
 3.  HASID....0027     PASID....0027     SASID..0027     JOBNAME..EZASOKGS
 4.  TCB......006E6A68 TIE......00008DF8 PLIST..00008E0C DUCB.....0000000C KEY..8
 5.  ADSNAME..GTASOKGS SUBTASK..MACROGIV                 TOKEN....7F6F3798 09902FB0
 6.  LOCAL  PORT..12035      IPADDR.. 9.67.113.58
 7.  REMOTE PORT..1034       IPADDR.. F901::32E1
 8.   REQAREA..:  00008D90x                                          Addr..00008D90
 9.   SOCKET...:  1                                                 Addr..00008A38
10.   NBYTE....:  40                                                Addr..00008A34
11.   ALET.....:  00000000x                                         Addr..000089A8
12.   BUF......:  (NO DATA)                                         Addr..000089A8
```

Line Description

1 This separator line shows the previous SYSTCPIP component trace record number in hexadecimal.

2 The first data line has the host name (MVS026), trace option (SOCKAPI), trace code (60050042), time, and trace record name.

3 The home, primary, and secondary ASIDs are always the same value (application's ASID) for the SOCKAPI trace option. The job name is also shown.

4 The MVS TCB address is shown. TIE (Task Interface Element) is the value of the TASK parameter on the EZASMI macro. The TIE is described in z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference. The parameter list address and DUCB are shown. Multiple concurrent calls can use the TIE; if so, they must have a different PLIST. The key is the 4-bit storage key from the PSW.

5 The ADSNAME (from the INITAPI call) is formatted in EBCDIC. The subtask name (from the INITAPI call) is formatted in EBCDIC if possible; otherwise, it is formatted in hexadecimal. The token is an 8-byte value, which identifies the INITAPI call instance.

6–7 If applicable, the ports and IP addresses are shown. The ports are formatted in decimal; the IP addresses are in dotted decimal.

8 The REQAREA parameter is shown because it was specified by the application. This is the 4-byte token presented to the application's exit when the response to the function request is complete. At the far right, the address in the application program of the REQAREA parameter is shown.

9 The SOCKET parameter is formatted in decimal. Its address is also shown.

10 The NBYTE parameter (number of bytes to be read) is formatted in decimal, followed by its address.

11 The ALET parameter is formatted in hexadecimal, followed by its address.

12 The BUF parameter currently has no data (because no data has been read) but its address is shown. In the READ Exit (or READ Async Complete) record, if the call was successful, the first 96 bytes of the data are also shown.

Examples of SOCKAPI trace records:
This topic includes descriptions and examples of the SOCKAPI trace records.

- "Successful API Call"
- "API call fails synchronously" on page 86
- "API call fails synchronously with parameter not addressable" on page 86
- "API call fails synchronously with diagnostic reason code" on page 86
- "Resolver API calls" on page 87
- "External IOCTL commands" on page 89
- "API call with an IOV parameter" on page 89
- "Default INITAPI" on page 90
- "Default TERMAPI" on page 90
- "SELECT" on page 90
- "SELECTEX" on page 91
- "Token Error" on page 91
- "Unsolicited event exit" on page 92

Successful API Call:
For asynchronous APIs, the Exit record merely indicates whether the call was acceptable. The contents of general-purpose register 15 are displayed to indicate this. The Asynchronous Complete record shows the actual results of the call. In addition to the output parameters, several interesting values are traced, including the contents of general-purpose register 0, the pointer to the asynchronous exit routine, the token passed to the asynchronous exit, the key in which the asynchronous exit was invoked, and the authorization state in which the exit is invoked. These values are not parameters on the GETHOSTID call, so their addresses are not shown. In this example, note also that the return code is formatted in dotted decimal and the meaning of the return code is provided.

Note: The API call might complete synchronously, in which case the Async Complete trace record might appear in the trace prior to the Exit record.

```
=================================================================00007B01
 MVS026    SOCKAPI   60050012  19:27:08.111729   GETHOSTID Exit
HASID....0027     PASID....0027     SASID..0027     JOBNAME..EZASOKOS
TCB......006E6A68 TIE......00006DF8 PLIST..00006E0C DUCB.....0000000C KEY..8
ADSNAME..EZASOKOS SUBTASK..00000000 00000000        TOKEN....7F6F3798 09902FB0
 REQAREA..:  00006D90x                                      Addr..00006D90
 R15......:  0 (CALL ACCEPTED)
=================================================================00007B05
 MVS026    SOCKAPI   60050032  19:27:08.111741   GETHOSTID Async Complete
HASID....0027     PASID....0027     SASID..0027     JOBNAME..EZASOKOS
TCB......006E6A68 TIE......00006DF8 PLIST..00006E0C DUCB.....0000000C KEY..8
ADSNAME..EZASOKOS SUBTASK..00000000 00000000        TOKEN....7F6F3798 09902FB0
 REQAREA..:  00006D90x                                      Addr..00006D90
 R0.......:  0x (NORMAL RETURN)
 ASYNC PTR:  00006B1C
 EXIT TOKEN: 00006B98x
```

```
 EXIT KEY.:  8x
 AUTHORIZATION STATE: PROBLEM
 RETCODE..:  9.67.113.58     (HOST IP ADDRESS)                  Addr..00006EB4
```

API call fails synchronously:
An asynchronous API call might fail synchronously or asynchronously. In this example, the WRITE call error was detected in the synchronous processing, so general-purpose register 15 has a nonzero value. The ERRNO value is interpreted (in this case, the NBYTE parameter on the WRITE call had a value of zero, which is not acceptable).

Note: The ERRNO value is the TCP/IP Sockets Extended Return Code. See z/OS Communications Server: IP and SNA Codes for information about TCP/IP Sockets Extended Return Codes.

```
=====================================================================00007B93
 MVS026    SOCKAPI   60050057  19:27:13.817195   WRITE Exit
HASID....0027     PASID....0027     SASID..0027     JOBNAME..EZASOKOS
TCB......006E6A68 TIE......00006DF8 PLIST..00006E0C DUCB.....00000009 KEY..8
ADSNAME..EZASOKOS SUBTASK..00000000 00000000        TOKEN....7F6F3798 09902FB0
LOCAL  PORT..11007      IPADDR.. 9.67.113.58
REMOTE PORT..1031       IPADDR.. 9.67.113.58
 REQAREA..:  00006D90x                                          Addr..00006D90
 SOCKET...:  1                                                  Addr..00006BDC
 R15......:  NON-ZERO (CALL WAS NOT ACCEPTED)
 ERRNO....:  10184 (EIBMWRITELENZERO)                           Addr..00006EB0
 RETCODE..:  -1                                                 Addr..00006EB4
```

API call fails synchronously with parameter not addressable:
If a parameter specified in the API call is not addressable by TCP/IP when creating the SOCKAPI record, the string (** PARAMETER NOT ADDRESSABLE **) is shown instead of the parameter value. The parameter address is shown at the far right, as usual.

```
=====================================================================00021347A
 VIC102    SOCKAPI   60050050  17:36:51.302111   SEND Entry

HASID....0026     PASID....0026     SASID..0026     JOBNAME..USER2
TCB......006D6D50 TIE......0000BDF8 PLIST..0000BE0C DUCB.....00000009 KEY..8
ADSNAME..USER2    SUBTASK..EZASOKEC                 TOKEN....7F755798 09806FB0
LOCAL  PORT..0          IPADDR ..0.0.0.0
REMOTE PORT..11007      IPADDR ..9.37.65.134
 SOCKET...:  0                                                  Addr..0000BA50
 NBYTE....:  96                                                 Addr..0000BA6C
 BUF......:  (** PARAMETER NOT ADDRESSABLE **)                  Addr..00015F38
 FLAGS....:  0 (NONE)                                           Addr..0000BC04
```

API call fails synchronously with diagnostic reason code:
If the API call does not complete successfully, the return code, ERRNO value (in decimal and interpreted), and possibly a diagnostic reason code are shown. The first two bytes of the diagnostic reason code are a qualifier (IBM internal use only). The last two bytes of the diagnostic reason codes are the UNIX ERRNOJR values described in z/OS Communications Server: IP and SNA Codes.

```
=====================================================================000085C1
 MVS026    SOCKAPI   60050004  19:36:01.934828   ACCEPT Exit
HASID....01F6     PASID....01F6     SASID..01F6     JOBNAME..EZASOKUE
TCB......006E6A68 TIE......00006DF0 PLIST..00006E04 DUCB.....0000000D KEY..8
ADSNAME..EZASOKUE SUBTASK..EZASOKUE                 TOKEN....7F6F3798 09902FB0
LOCAL  PORT..11007      IPADDR ..9.67.113.58
REMOTE PORT..0          IPADDR ..0.0.0.0
 REQAREA..:  00000000x                                          Addr..00006D80
 SOCKET...:  0                                                  Addr..00006BA8
 NAME.....:  (NO DATA)                                          Addr..00006BAC
```

```
 DIAG. RSN:  76620291x
 ERRNO....:  5 (EIO)                                              Addr..00006EA8
 RETCODE..:  -1                                                   Addr..00006EAC
```

Resolver API calls:
The GETHOSTBYADDR and GETHOSTBYNAME IPv4 Resolver API calls use the HOSTENT structure described in the calls in z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference. As shown in the following GETHOSTBYADDR Exit trace example, the HOSTENT address is shown on one line, and the contents of the HOSTENT structure are described on separate lines. There can be multiple aliases and host addresses; each one is listed separately. In this example, there are two aliases.

```
=================================================================000051CB
 MVS026    SOCKAPI   60050066  19:02:01.426345   GETHOSTBYADDR Exit

HASID....0027     PASID....0027     SASID..0027     JOBNAME..EZASOKGH
TCB......006E6A68 TIE......00007DF8 PLIST..00007E0C DUCB.....0000000A KEY..0
ADSNAME..EZASOKGH SUBTASK..00000000 00000000        TOKEN....00000000 09902FB0
 HOSTENT..:                                                       Addr..00005F08
 HOSTNAME.:                                                       Addr..00005F30
 Loopback
 FAMILY...:  2                                                    Addr..00005F10
 ADDR LEN.:  4                                                    Addr..00005F14
 HOSTADDR.:  127.0.0.1                                            Addr..00005F54
 ALIAS....:  LOOPBACK                                             Addr..00005F3C
 ALIAS....:  LOCALHOST                                            Addr..00005F48
 RETCODE..:  0                                                    Addr..00007EB4
```

The GETADDRINFO for IPv4 or IPv6 Resolver API shows that the call is requesting the IP address for the host (node) name MVS150. No service name is provided. GETADDRINFO exit shows that the hostname was resolved to the IPv4 address 9.67.113.117. These fields are described in the Macro and CALL section in z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

```
=================================================================0000134C
  MVS150    SOCKAPI   6005006D  15:06:07.294268   GETADDRINFO Entry
 HASID....002D     PASID....002D     SASID..002D     JOBNAME..USER1X
 TCB......007F63B0 TIE......0A90AAD8 PLIST..0A90AAEC DUCB.....00000009 KEY..8
 ADSNAME.......... SUBTASK..EZASO6CS                TOKEN....7F694220 0A97EFB0
  NODELEN..:  6                                                   Addr..0A973490
  NODE.....:                                                      Addr..0A973390
       MVS150
  SERVLEN..:  0                                                   Addr..0A9734B8
  SERVICE..:  (NO DATA)                                           Addr..0A973498
  HINTS....:  0A913F70x (ADDRINFO Address)                        Addr..0A913F90
  ADDRINFO Structure..:
   AF.......  0 (AF_UNSPEC)     FLAGS..... 00000002x
   SOCTYPE..  0 (UNKNOWN)       PROTO..... 0 (IPPROTO_IP)
   NAME.....  00000000x         NAMELEN... 0
   CANONNAME  00000000x         NEXT...... 00000000x
  CANNLEN..:  (NO DATA)                                           Addr..0A9734C0
  RES......:  (NO DATA)                                           Addr..0A913F94
 =================================================================0000134D
  MVS150    SOCKAPI   6005006E  15:06:09.997756   GETADDRINFO Exit
 HASID....002D     PASID....002D     SASID..002D     JOBNAME..USER1X
 TCB......007F63B0 TIE......0A90AAD8 PLIST..0A90AAEC DUCB.....00000009 KEY..8
 ADSNAME.......... SUBTASK..EZASO6CS                TOKEN....7F694220 0A97EFB0
  HINTS....:  0A913F70x (ADDRINFO Address)                        Addr..0A913F90
  ADDRINFO Structure..:
   AF.......  0 (AF_UNSPEC)     FLAGS..... 00000002x
   SOCTYPE..  0 (UNKNOWN)       PROTO..... 0 (IPPROTO_IP)
   NAME.....  0002111Cx         NAMELEN... 0
    PORT....  0                 IPADDR.... 0.0.0.0
```

```
   FAMILY..  0 (UNKNOWN)       RESERVED.. 0000000000000000x
  CANONNAME  00000000x         NEXT...... 00000000x
 CANNLEN..:  22                                              Addr..0A9734C0
 RES......:  0002111Cx (ADDRINFO Address)                    Addr..0A913F94
 ADDRINFO Structure..:
  AF.......  2 (AF_INET)       FLAGS..... 00000000x
  SOCTYPE..  1 (STREAM)        PROTO..... 0 (IPPROTO_IP)
  NAME.....  0002114Cx         NAMELEN... 16
   PORT....  0                 IPADDR.... 9.67.113.117
   FAMILY..  2 (AF_INET)       RESERVED.. 0000000000000000x
  CANONNAME  0002101Cx         NEXT...... 00000000x
  MVS150.raleigh.ibm.com
```

The FREEADDRINFO for IPv4 or IPv6 Resolver API call displays the RES (ADDRINFO) structure that is freed. This field is in the Macro and CALL section in z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

```
=================================================================0000134E
 MVS150    SOCKAPI   6005006F  15:06:09.998002   FREEADDRINFO Entry
HASID....002D     PASID....002D     SASID..002D     JOBNAME..USER1X
TCB......007F63B0 TIE......0A90AAD8 PLIST..0A90AAEC DUCB.....00000009 KEY..8
ADSNAME.......... SUBTASK..EZASO6CS                 TOKEN....7F694220 0A97EFB0
 ADDRINFO.:  0002111Cx (ADDRINFO Address)                    Addr..0A913F94
 ADDRINFO Structure..:
  AF.......  2 (AF_INET)       FLAGS..... 00000000x
  SOCTYPE..  1 (STREAM)        PROTO..... 0 (IPPROTO_IP)
  NAME.....  0002114Cx         NAMELEN... 16
   PORT....  0                 IPADDR.... 9.67.113.117
   FAMILY..  2 (AF_INET)       RESERVED.. 0000000000000000x
  CANONNAME  0002101Cx         NEXT...... 00000000x
  MVS150.raleigh.ibm.com
=================================================================0000134F
 MVS150    SOCKAPI   60050070  15:06:09.999021   FREEADDRINFO Exit
HASID....002D     PASID....002D     SASID..002D     JOBNAME..USER1X
TCB......007F63B0 TIE......0A90AAD8 PLIST..0A90AAEC DUCB.....00000009 KEY..8
ADSNAME.......... SUBTASK..EZASO6CS                 TOKEN....7F694220 0A97EFB0
```

The GETNAMEINFO for IPv4 or IPv6 Resolver API shows that the call is requesting the name of the IPv6 address ::1 and the service name for port 1031. GETNAMEINFO Exit shows that the IP address was resolved to the name loop6int.resdns.ibm.com and no service name was found for port 1031 (hence the service name is the input port number). These fields are in the Macro and CALL section in z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

```
=================================================================0000135F
 MVS150    SOCKAPI   6005006B  15:06:45.481639   GETNAMEINFO Entry
HASID....0025     PASID....0025     SASID..0025     JOBNAME..USER1Y
TCB......007F62F8 TIE......0A903AD8 PLIST..0A903AEC DUCB.....00000008 KEY..8
ADSNAME.......... SUBTASK..EZSO6CC                  TOKEN....7F6E2220 0A977FB0
 NAMELEN..:  28                                              Addr..0A96C348
 NAME.....:                                                  Addr..0A96C500
   PORT... 1031          IPADDR.. ::1
   FAMILY. 19 (AF_INET6) LENGTH.. 0
   FLOWINFO. 00000000x   SCOPID.. 00000000x
 HOSTLEN..:  255                                             Addr..0A96C450
 HOST.....:  (NO DATA)                                       Addr..0A96C350
 SERVLEN..:  32                                              Addr..0A96C478
 SERVICE..:  (NO DATA)                                       Addr..0A96C458
 FLAGS....:  00000004x                                       Addr..0A96C480
=================================================================00001360
 MVS150    SOCKAPI   6005006C  15:06:46.707053   GETNAMEINFO Exit
HASID....0025     PASID....0025     SASID..0025     JOBNAME..USER1Y
TCB......007F62F8 TIE......0A903AD8 PLIST..0A903AEC DUCB.....00000008 KEY..8
ADSNAME.......... SUBTASK..EZSO6CC                  TOKEN....7F6E2220 0A977FB0
```

```
  HOSTLEN..:  23                                                  Addr..0A96C450
  HOST.....:                                                      Addr..0A96C350
       loop6int.resdns.ibm.com
  SERVLEN..:  4                                                   Addr..0A96C478
  SERVICE..:  1031                                                Addr..0A96C458
  FLAGS....:  00000004x                                           Addr..0A96C480
```

External IOCTL commands:
For external IOCTL commands, the command name is interpreted. For IBM internal-use-only commands, the hexadecimal value of the command is shown. The input and output for each command can differ. In this example, the SIOCGIFCONF command requests the network interface configuration. The exit record shows the call was successful (the return code is zero) and the network interface configuration is shown.

```
=====================================================================00001734
  MVS026    SOCKAPI   6005001F  20:42:44.805938   IOCTL Entry

 HASID....19       PASID....19       SASID..19       JOBNAME..USER1
 TCB......006AFD40 TIE......68DF8    PLIST..00068E0C DUCB.....00000008 KEY..8
 ADSNAME..USER1    SUBTASK..00000000 00000000        TOKEN....7F67F798 0A2B4FB0
 LOCAL  PORT..11007     IPADDR ..9.67.113.58
 REMOTE PORT..0         IPADDR ..0.0.0.0
  SOCKET...:  0                                                  Addr..000685A0
  COMMAND..:  SIOCGIFCONF                                        Addr..0006782C
  REQARG...:                                                     Addr..00068928
    BUFFER LENGTH.. 99
 =====================================================================00000323
  MVS026    SOCKAPI   60050020  20:42:44.806101   IOCTL Exit

 HASID....19       PASID....19       SASID..19       JOBNAME..USER1
 TCB......006AFD40 TIE......68DF8    PLIST..00068E0C DUCB.....00000008 KEY..8
 ADSNAME..USER1    SUBTASK..00000000 00000000        TOKEN....7F67F798 0A2B4FB0
 LOCAL  PORT..11007     IPADDR ...9.67.113.58
 REMOTE PORT..0         IPADDR ...0.0.0.0
  SOCKET...:  0                                                  Addr..000685A0
  COMMAND..:  SIOCGIFCONF                                        Addr..0006782C
  RETARG...:                                                     Addr..000685C4
  Socket Name.. TR1
    PORT....  0                IPADDR.... 9.67.113.58
    FAMILY..  2 (AF_INET)      RESERVED.. 0000000000000000x
  RETCODE..:  0                                                  Addr..00068EB4
```

API call with an IOV parameter:
The IOV parameter is an array of structures used on the READV, RECVMSG, SENDMSG, and WRITEV API calls. Each structure contains three words: the buffer address, the ALET, and the buffer length. Each IOV entry is shown on one line. When there is data available (READV Exit, RECVMSG Exit, SENDMSG Entry, and WRITEV Entry), some of the buffer data is also displayed. A maximum of 96 bytes of data are displayed.

In the READV Exit example, three IOV entries were specified, but only two were used. All the data is displayed because the total is less than 96 bytes.

```
=====================================================================00001773
 MVS026    SOCKAPI   60050045  19:19:20.954789   READV Exit

HASID....0024     PASID....0024     SASID..0024     JOBNAME..EZASOKKS
TCB......006E6A68 TIE......00007DF8 PLIST..00007E0C DUCB.....0000000B KEY..8
ADSNAME..EZASOKKS SUBTASK..EZASOKKS                 TOKEN....7F6F3798 09902FB0
LOCAL  PORT..11007     IPADDR ..9.67.113.58
REMOTE PORT..1032      IPADDR ..9.67.113.58
 REQAREA..:  00007D90x                                          Addr..00007D90
 SOCKET...:  1                                                  Addr..0000776C
 IOVCNT...:  3                                                  Addr..000077B4
```

```
 IOVENTRY.:  LENGTH..10     ALET..0x                            Addr..00007890
     +0000  E38889A2  4089A240  8396               | This is co       |
 IOVENTRY.:  LENGTH..10     ALET..0x                            Addr..0000789A
     +0000  99998583  A34B                         | rrect.           |
 IOVENTRY.:  LENGTH..10     ALET..0x                            Addr..000078A4
 RETCODE..:  16 BYTES TRANSFERRED                               Addr..00007EB4
```

Default INITAPI:
An explicit INITAPI call is not required prior to some API calls, so TCP/IP creates a default INITAPI. (See z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference for the complete list.) The default INITAPI record is traced after the Entry record for the API call that caused the default INITAPI to occur. There is just one record for this event (no Exit record).

```
=================================================================000077EC
 MVS026    SOCKAPI   60050040  19:24:11.552924   Default INITAPI

HASID....0027     PASID....0027     SASID..0027     JOBNAME..EZASOKSX
TCB......006E6A68 TIE......00007DF0 PLIST..00007E04 DUCB.....0000000A KEY..8
ADSNAME..EZASOKSX SUBTASK..00000000 00000000        TOKEN....7F6F3798 09902FB0
 MAXSNO...:  49
 APITYPE..:  2
 RETCODE..:  0
```

Default TERMAPI:
Usually, an application ends the connection between itself and TCP/IP by issuing the TERMAPI call. But sometimes, the connection ends for another reason, such as the application being canceled. In this case, TCP/IP issues a default TERMAPI. The default TERMAPI is traced in a SOCKAPI trace record. There is just one record for this event (no Exit record).

```
=================================================================00000168
  MVS026    SOCKAPI   60050069  22:46:48.185419   Default TERMAPI

 HASID....01F9     PASID....01F9     SASID..01F9     JOBNAME..EZASOKQS
 TCB......006E6A68 TIE......08920888 PLIST..00000000 DUCB.....00000008 KEY..6
 ADSNAME..EZASOKQS SUBTASK..EZASOKQS                 TOKEN....7F6F3798 00000000
```

SELECT:
For SELECT and SELECTEX, the socket masks are formatted in both binary and decimal. The socket list is displayed first in binary. The socket numbers are indicated by the bit position in the mask, starting with bit position 0 (for socket 0), which is the rightmost bit. The bit positions (socket numbers) are shown at left.

For example, the lowest numbered sockets are on the last line; they are sockets 0 - 31. In this line, only sockets 0, 1, 2, and 3 are selected. As shown in the following example, the binary mask, the decimal socket numbers are listed in numerical order. This is a convenient way to check whether the mask is coded as intended.

```
=================================================================00024EDF
  BOTSWANA  SOCKAPI   6005004C  20:51:35.477605   SELECT Entry

HASID....0078     PASID....0078     SASID..0078     JOBNAME..TN1
TCB......007F6988 TIE......1463227C PLIST..1477EF18 DUCB.....00000016 KEY..8
ADSNAME..         SUBTASK..14632138                 TOKEN....7F75FFC8 1468FA90
 REQAREA..:  1477EEF0x                                        Addr..1477EF98
 MAXSOC...:  100                                              Addr..14632258
 TIMEOUT..:  SECOND..0      MICRO SECOND..500000              Addr..1463226C
 RSNDMSK..:                                                   Addr..14632108
 SOCKET NO.    READ SOCKET MASK (INPUT)
 (Decimal)     (Binary)
 31  0  00000000 00000000 00000000 00001111
 63  32 00111011 11111111 11111111 11111101
 95  64 11111111 11111111 10111111 11111111
```

```
127 96 00000000 00000000 00000000 11110111
  SELECTED SOCKETS:
   0,     1,     2,     3,    32,    34,    35,    36,    37,    38
  39,    40,    41,    42,    43,    44,    45,    46,    47,    48
  49,    50,    51,    52,    53,    54,    55,    56,    57,    59
  60,    61,    64,    65,    66,    67,    68,    69,    70,    71
  72,    73,    74,    75,    76,    77,    79,    80,    81,    82
  83,    84,    85,    86,    87,    88,    89,    90,    91,    92
  93,    94,    95,    96,    97,    98,   100,   101,   102,   103
```

If the MAXSOC value is so large that all the SELECT or SELECTEX parameters cannot be traced within a single 14K buffer, multiple trace entries are written (one trace entry for each mask). When multiple trace entries are written for the same SELECT or SELECTEX call entry or exit, all the trace data except the masks themselves are duplicated across the trace entries. For example, the time stamp is the same, the MAXSOC value is the same, the TIMEOUT value is the same, and so on. The trace description indicates to which mask the trace entries pertain. For example, if the MAXSOC value in the above trace example were 65535, then each mask would be traced individually.

```
=====================================================================00024EDF
  BOTSWANA  SOCKAPI   6005004C  20:51:35.477605   SELECT Entry (read mask)

HASID....0078     PASID....0078     SASID..0078     JOBNAME..TN1
TCB......007F6988 TIE......1463227C PLIST..1477EF18 DUCB.....00000016 KEY..8
ADSNAME..         SUBTASK..14632138                 TOKEN....7F75FFC8 1468FA90
 REQAREA..:  1477EEF0x                                          Addr..1477EF98
 MAXSOC...:  65535                                              Addr..14632258
 TIMEOUT..:  SECOND..0      MICRO SECOND..500000                Addr..1463226C
 RSNDMSK..:                                                     Addr..14632108
 SOCKET NO.    READ SOCKET MASK (INPUT)
 (Decimal)     (Binary)
```

SELECTEX:
The SELECTEX call can contain a list of ECBs. The high-order bit on the SELECB address indicates whether a list of ECBs was specified. Since the high-order bit is on in this example, there is a list of ECBs. The end of the list is indicated by the high-order bit in the ECB address. In this example, the time limit expired before any ECBs were posted. Since no selected sockets were ready, the read, write, and error masks indicate that there is no data to report.

```
=====================================================================000078FB
 MVS026    SOCKAPI   6005004F  19:25:48.610379   SELECTEX Exit

HASID....0027     PASID....0027     SASID..0027     JOBNAME..EZASOKX4
TCB......006E6A68 TIE......00007DF8 PLIST..00007E0C DUCB.....0000000C KEY..8
ADSNAME..EZASOKX4 SUBTASK..BARBARA                  TOKEN....7F6F3798 09902FB0
 MAXSOC...:  33                                                 Addr..00007AE8
 TIMEOUT..:  SECOND..0      MICRO SECOND..35                    Addr..00007AF4
 RRETMSK..:  (NO DATA)                                          Addr..00007B0C
 WRETMSK..:  (NO DATA)                                          Addr..00007B14
 ERETMSK..:  (NO DATA)                                          Addr..00007B1C
 SELECB...:                                                     Addr..80007B60
 ECB......:  00000000x                                          Addr..00007B70
 ECB......:  00000000x                                          Addr..00007B74
 ECB......:  00000000x                                          Addr..00007B78
 ECB......:  00000000x                                          Addr..80007B7C
 RETCODE..:  0 (TIME LIMIT EXPIRED)                             Addr..00007EB4
```

Token Error:
When an API call fails early in processing, before the SOCKAPI Entry record is created, the Token Error SOCKAPI record is written. In the example, the BIND call failed due to the token being overwritten. (The token at offset eight has X'FFFF'.) There is no BIND Entry or Exit record.

```
=================================================================00000158
 MVS026    SOCKAPI   6005006A  22:46:48.173348   Token Error

HASID....01F9     PASID....01F9     SASID..01F9     JOBNAME..EZASOKQS
TCB......006E6A68 TIE......00006DF8 PLIST..00006E0C DUCB.....00000008 KEY..8
ADSNAME..         SUBTASK..                         TOKEN....7F6F3798 09902FB0
 CALL.....:  BIND
 TOKEN....:  7F6F3798 09902FB0 FFFF0000 00003FC5x
 ERRNO....:  1028 (EIBMINVTCPCONNECTION)
 RETCODE..:  -1
```

Unsolicited event exit:
If the unsolicited event exit is driven, a SOCKAPI trace record is created (if the SOCKAPI trace option is active).

Note: The key in the header is 0. This means that the UEE trace record was created when TCP/IP was in key zero. The UEEXIT has key 8, which means the UE exit is invoked in key eight.

```
=================================================================000086FC
 MVS026    SOCKAPI   60050041  19:36:04.965468   Unsolicited Event Exit Invoked

HASID....0024     PASID....0024     SASID..0024     JOBNAME..TCPIPROC
TCB......006E6A40 TIE......00006DF0 PLIST..00000000 DUCB.....00000000 KEY..0
ADSNAME..EZASOKUE SUBTASK..EZASOKUE                 TOKEN....7F6F3798 00000000
 UEEXIT...:  ADDRESS..00006B30  TOKEN..00006D80x   ASCB.....00F94C80x KEY..8
             REASON...1 (TCP/IP TERMINATION)
```

Correlating the data trace and packet trace with the SOCKAPI trace

The SOCKAPI option only records the first 96 bytes of data. To see all the data that was sent or received, you must also activate the data trace or packet trace. The data trace can be correlated easily with the SOCKAPI trace option because both traces are recording data between the application and the TCP/IP stack. The traces can be merged with the IPCS MERGE subcommand. The data trace header contains fields that allow the full data to be correlated.

Figure 16 on page 93 shows the data trace record corresponding to the READ Exit SOCKAPI trace entry in Figure 17 on page 93. The server issues READ and waits for a message. The data trace record shows the entire 120 bytes of data because the FULL option was used when starting the data trace. In the READ Exit record, only the first 96 bytes of data are shown.

The records in the two traces can be correlated by the following:

Time The data trace time must be prior to the READ Exit record time. The data trace time is `20:08:09.181239`. The READ Exit record time is `20:08:09.181354`.

Jobname
The job name is EZASOKAS in both records.

ASID The ASID is the server's 0024 (hexadecimal) in both records.

TCB The TCB is `006E6A68` in both records.

Data length
In the data trace, the length is 78 hexadecimal, which is 120 decimal. The SOCKAPI trace record shows that the return code is 120 (decimal) bytes.

Port The source port number in the data trace record (11007 decimal) matches the local port number in the SOCKAPI trace record. The destination and remote ports also match (1040 decimal).

IP Address

The IP addresses are handled in the same way as the port numbers. In this example, both the client and server were on the same TCP/IP stack, so the IP addresses are the same.

```
 MVS026    DATA      00000003  20:08:09.181239  Data Trace

 JOBNAME =    EZASOKAS          FROM  FULL
  TOD CLOCK = XB395B2C2  40035C03
  PKT 2           LOST RECORDS = 0      HDR SEQUENCE NUM = 1
 SOURCE IP ADDR = 9.67.113.58      DEST IP ADDR = 9.67.113.58
 SOURCE PORT = 11007  DEST PORT = 1040   ASID = X0024  TCB = X006E6A68
  DATA LENGTH = X0078
0000 E38889A2 4089A240 8140A2A3 99899587 *This is a string|....@..@.@......*
0010 40A689A3 88408696 99A3A840 83888199 * with forty char|@....@.....@....*
0020 8183A385 99A24B40 E38889A2 4089A240 *acters. This is |......K@....@..@*
0030 8140A2A3 99899587 40A689A3 88408696 *a string with fo|.@......@....@..*
0040 99A3A840 83888199 8183A385 99A24B40 *rty characters.  |...@..........K@*
0050 E38889A2 4089A240 8140A2A3 99899587 *This is a string|....@..@.@......*
0060 40A689A3 88408696 99A3A840 83888199 * with forty char|@....@.....@....*
0070 8183A385 99A24B40                   *acters.         |......K@        *
```

Figure 16. Data trace record.

```
=====================================================================00002403
 MVS026    SOCKAPI   60050043  20:08:09.181354   READ Exit

HASID....0024     PASID....0024     SASID..0024     JOBNAME..EZASOKAS
TCB......006E6A68 TIE......00006DF8 PLIST..00006E0C DUCB.....00000009 KEY..8
ADSNAME..EZASOKAS SUBTASK..EZASOKAS                 TOKEN....7F6F3798 09902FB0
LOCAL  PORT..11007     IPADDR ..9.67.113.58
REMOTE PORT..1040      IPADDR ..9.67.113.58
REQAREA..:  00006D90x                                          Addr..00006D90
SOCKET...:  1                                                  Addr..00006B94
NBYTE....:  120                                                Addr..00006B90
BUF......:                                                     Addr..00006B96
    +0000  E38889A2  4089A240  8140A2A3  99899587  | This is a string |
    +0010  40A689A3  88408696  99A3A840  83888199  |  with forty char |
    +0020  8183A385  99A24B40  E38889A2  4089A240  | acters. This is  |
    +0030  8140A2A3  99899587  40A689A3  88408696  | a string with fo |
    +0040  99A3A840  83888199  8183A385  99A24B40  | rty characters.  |
    +0050  E38889A2  4089A240  8140A2A3  99899587  | This is a string |
RETCODE..:  120 BYTES TRANSFERRED                              Addr..00006EB4
=====================================================================00002407
```

Figure 17. SOCKAPI trace record.

The packet trace, on the other hand, does not correlate well with the SOCKAPI trace option. The packet trace records data being sent or received between the TCP/IP stack and the network. The packet trace data has headers and the data can be segmented or packed.

Packet trace (SYSTCPDA) for TCP/IP stacks

Packet trace is a diagnostic method for obtaining traces of IP packets flowing to and from a TCP/IP stack on a z/OS Communications Server host. You can use the PKTTRACE statement to copy IP packets as they enter or leave TCP/IP, and then examine the contents of the copied packets. To be traced, an IP packet must meet all the conditions specified on the PKTTRACE statement.

The trace process

Trace data is collected as IP packets enter or leave TCP/IP. The actual collection occurs within the device drivers of TCP/IP, which capture the data that is received from or sent to the network.

Packets that are captured have extra information added to them before they are stored. This extra information is used during the formatting of the packets. The captured data reflects exactly what the network sees. For example, the trace contains the constituent packets of a fragmented packet exactly as they are received or sent.

The selection criteria for choosing packets to trace are specified through the PKTTRACE statement for the TCP/IP address space. See z/OS Communications Server: IP System Administrator's Commands for more information about the PKTTRACE statement and subcommand.

The PKTTRACE statement and subcommand are applied to device links that are defined in the TCP/IP address space through the LINK statement. Figure 18 illustrates the overall control and data flow in the IP packet tracing facility.

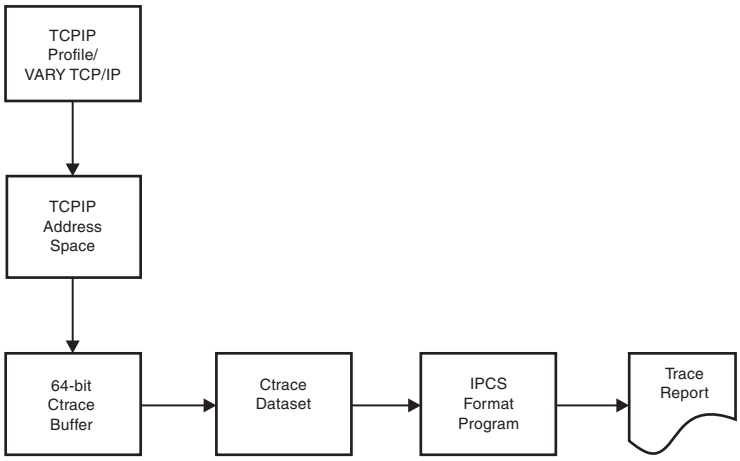

Figure 18. Control and data flow in the IP packet tracing facility

Supported devices

IP packet tracing is supported for all network interfaces supported by TCP/IP (including loopback).

When using the MULTIPATH option of the IPCONFIG statement, packets can be sent over multiple interfaces. All of the interfaces must be traced. In this case specify an IP address to select the required packet. This statement also applies to the case where packets can be received over multiple interfaces (even if MULTIPATH is not used by this TCP/IP).

For information about the format of the packet trace command (VARY PKTTRACE) see z/OS Communications Server: IP System Administrator's Commands.

Starting packet trace

To start packet trace, use the following command:

```
V TCPIP,tcpprocname,PKT
```

Security Note: To use any VARY command, the user must be authorized in RACF.

The RACF profile for each user must have access for a resource of the form MVS.VARY.TCPIP.*xxx*, where *xxx* is the first eight characters of the command name. For packet trace, this would be MVS.VARY.TCPIP.PKTTRACE.

Traces are placed in an internal buffer, which can then be written out using an external writer. The MVS TRACE command must also be issued for component SYSTCPDA to activate the packet trace.

After starting packet trace, you can display the status using the Netstat command, as shown in the following example:

```
NETSTAT -p TCPCS -d
MVS TCP/IP NETSTAT CS V1R10 TCPIP Name: TCPCS 18:03:31
DevName: LOOPBACK          DevType: LOOPBACK
DevStatus: Ready
LnkName: LOOPBACK          LnkType: LOOPBACK   LnkStatus: Ready
  NetNum:  0   QueSize: 0
  BytesIn: 192537                  BytesOut: 192537
  ActMtu:  65535
BSD Routing Parameters:
  MTU Size: 00000             Metric: 00
  DestAddr: 0.0.0.0           SubnetMask: 0.0.0.0
Packet Trace Setting:
  Protocol: *                 TrRecCnt: 00000000  PckLength: FULL
  Discard :  0
  SrcPort:  *                 DestPort: *         PortNum *
  IpAddr:   9.67.113.1        SubNet:   *
Multicast Specific:
  Multicast Capability: No
```

In this example, the packet length (PckLength) is FULL and TrRecCnt is the number of packets written for this device.

Note: If you are a TSO user, use the corresponding NETSTAT DEV command.

Modifying options with VARY

After starting a packet trace, you can change the trace using the VARY command. For example, if you want to change the packet trace to abbreviate the data being traced, use the following command:

```
V TCPIP,tcpproc,PKT,ABBREV
```

You can display the results of the VARY command using the Netstat command:

```
NETSTAT -p TCPCS -d
MVS TCP/IP NETSTAT CS V1R10 TCPIP Name: TCPCS 18:03:31
DevName: LOOPBACK          DevType: LOOPBACK
  DevStatus: Ready
  LnkName: LOOPBACK          LnkType: LOOPBACK   LnkStatus: Ready
    NetNum:  0   QueSize: 0
    BytesIn: 813                     BytesOut: 813
    ActMtu:  65535
  BSD Routing Parameters:
    MTU Size: 00000             Metric: 00
    DestAddr: 0.0.0.0           SubnetMask: 0.0.0.0
  Packet Trace Setting:
```

```
    Protocol: *                 TrRecCnt: 00000000  PckLength: 00200
    Discard:  0
    SrcPort:  *                 DestPort: *         PortNum: *
    IpAddr:   *                 SubNet:   *
  Multicast Specific:
    Multicast Capability: No
```

Tip: If you are a TSO user, use the corresponding NETSTAT option.

By issuing multiple VARY commands, you can OR filters together. For example, issuing the following VARY commands records all packets whose destination port is *xxxx* or whose source port is *xxxx*.

```
V TCPIP,tcpproc,PKTTRACE,DEST=xxxx
V TCPIP,tcpproc,PKTTRACE,SRCP=xxxx
```

The result is a trace that contains only packets with a source port of *xxxx* or packets with a destination port of *xxxx*.

Tip: An alternative command to use is the PKTTRACE command with PORTNUM.

```
V TCPIP,,PKTTRACE,PORTNUM=xxxx
```

If both DEST and SRCP are specified in the same command, you can AND the parameters together. For example, issuing the following VARY command records only the packets with both a destination port of *xxxx* and a source port of *yyyy*.

```
V TCPIP,tcpproc,PKTTRACE,DEST=xxxx,SRCP=yyyy
```

You can use the VARY TCPIP,*tcpproc*,OBEYFILE command to make temporary dynamic changes to system operation and network configuration without stopping and restarting the TCP/IP address space. For example, if you started the address space TCPIPA and created a sequential data set USER99.TCPIP.OBEYFIL1 containing packet trace statements, issue the following command:

```
VARY TCPIP,TCPIPA,CMD=OBEYFILE,DSN=USER99.TCPIP.OBEYFIL1
```

The VARY TCPIP,PKTTRACE command is cumulative. You can trace all packets for specified IP addresses by entering multiple PKTTRACE commands. In the following example, the two commands trace all the packets received and all the packets sent for the specified IP addresses.

```
VARY TCPIP,,PKT,ON,IPADDR=10.27.142.44
VARY TCPIP,,PKT,ON,IPADDR=10.27.142.45
```

Formatting packet traces using IPCS

The IPCS CTRACE command parameters are described in "Formatting component traces" on page 53. The following notes apply to the IPCS CTRACE parameters about the packet trace formatter:

JOBLIST, JOBNAME
The LINKNAME and JOBNAME keywords in the OPTIONS string can also be used to select records.

TALLY
Equivalent to the STATISTICS(DETAIL) option.

START and STOP
Packets are numbered after the START keyword has filtered records.

LIMIT
See the RECORDS keyword in the OPTIONS string.

USEREXIT
The packet trace formatter calls the CTRACE USEREXIT before testing the records with the filtering criteria. If it returns a nonzero return code, then the record is skipped. The USEREXIT can also be used in the OPTIONS string. It is called after the record has met all the filtering criteria in the OPTIONS string.

COMP
Must be SYSTCPDA.

SUB
The SUB must name the TCP/IP procedure that created the CTRACE records when the input is a dump data set.

EXCEPTION
Since there are no EXCEPTION records for packet trace, the EXCEPTION keyword must not be specified.

ENTIDLIST
The following are the valid values for packet trace:

1 IPv4 packet trace records

2 X25 trace records

3 IPv4 Enterprise Extender data trace records

Tip: Type 1, Type 2, and Type 3 records are no longer written by TCP/IP.

4 IPv4 and IPv6 packet trace records

5 IPv4 and IPv6 data trace records

6 Enterprise Extender trace records

The CTRACE OPTIONS string provides a means of entering additional keywords for record selection and formatting packet traces (COMP=SYSTCPDA). See "Syntax" on page 55 for the complete syntax of CTRACE.

OPTIONS syntax

OPTIONS component

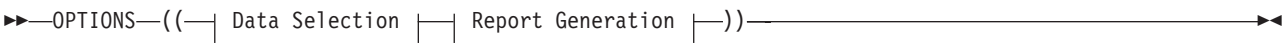

Data Selection:

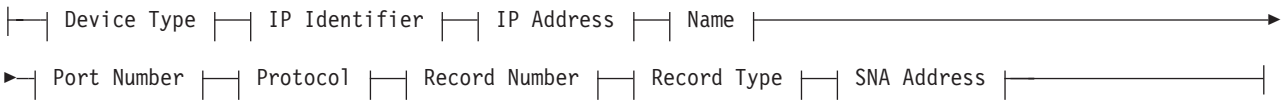

Device Type:

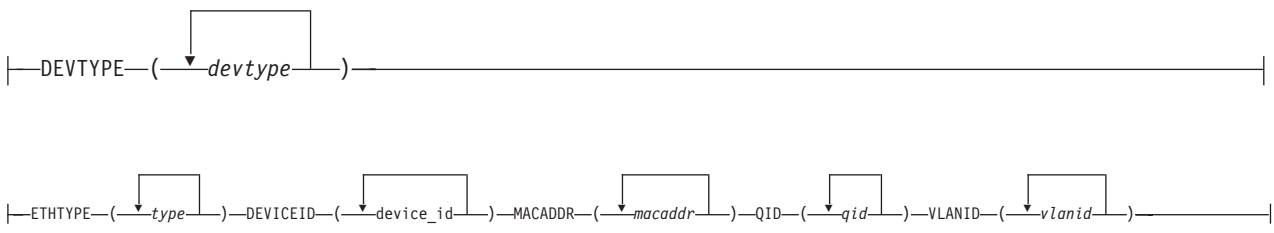

IP Identifier:

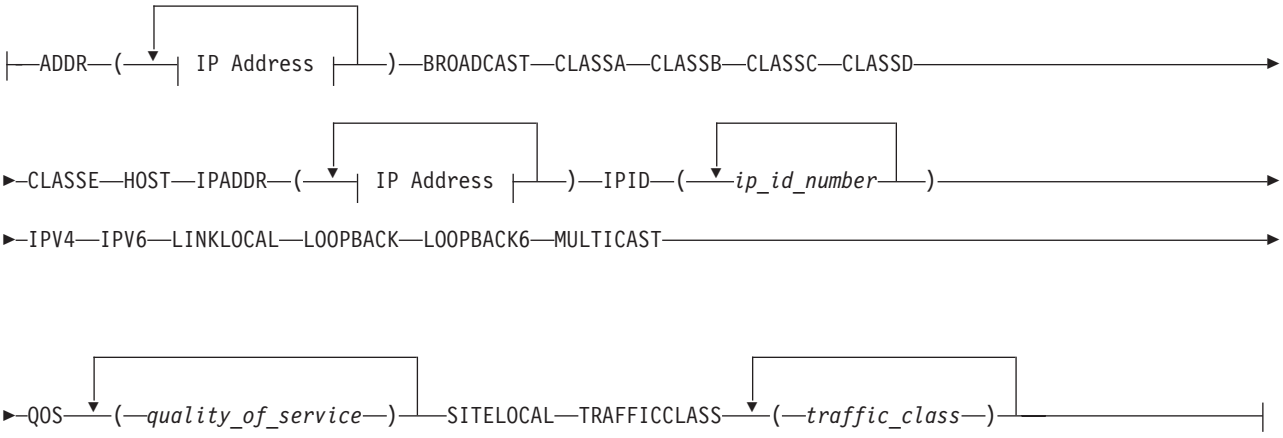

IP Address:

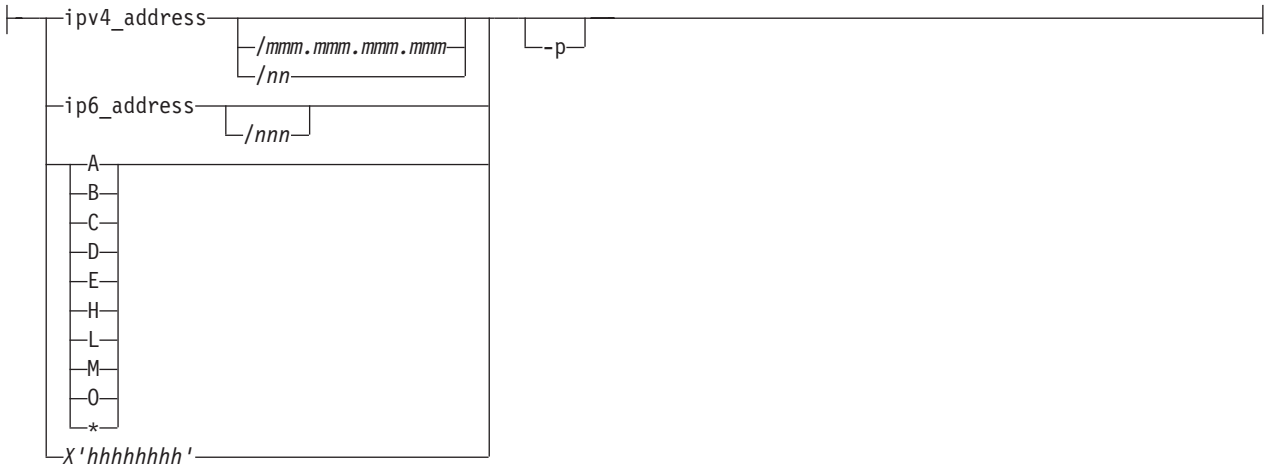

Name:

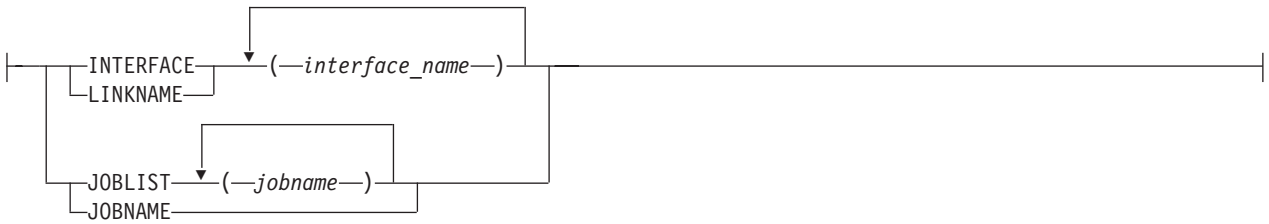

Port Number:

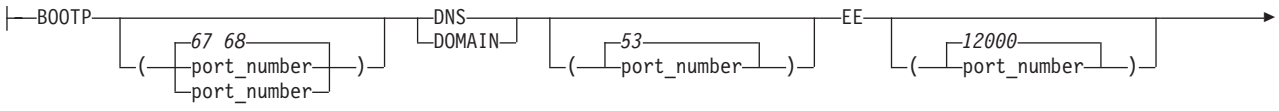

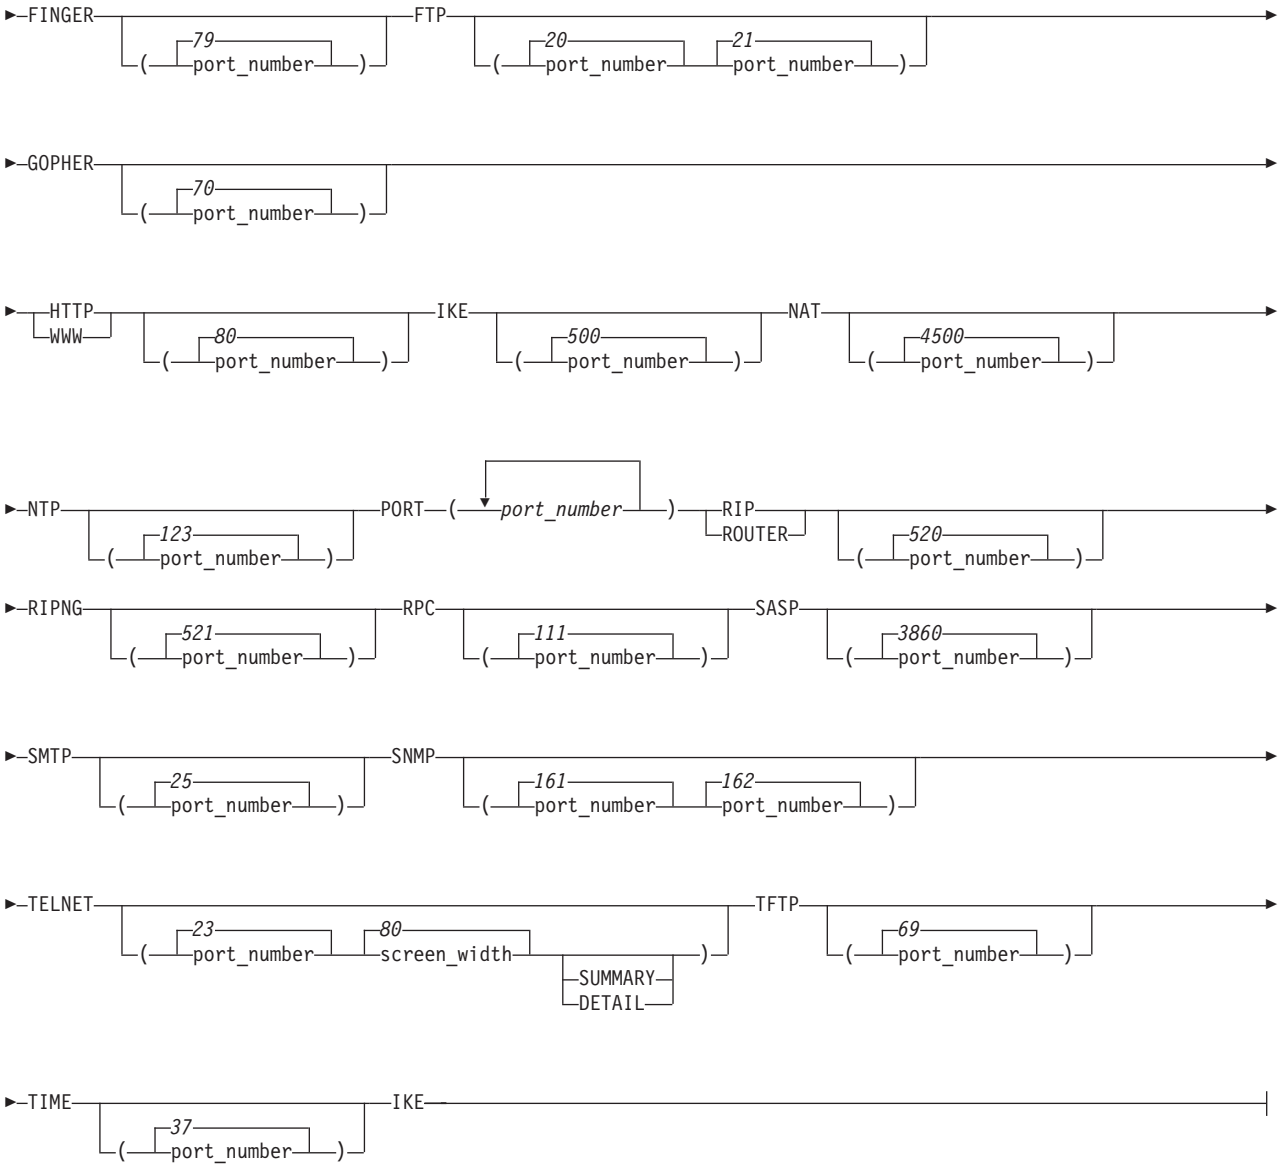

Protocol:

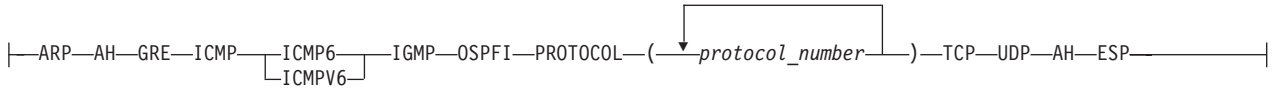

Record Number:

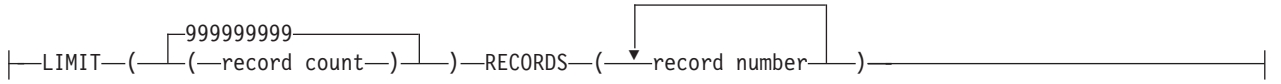

Record Type:

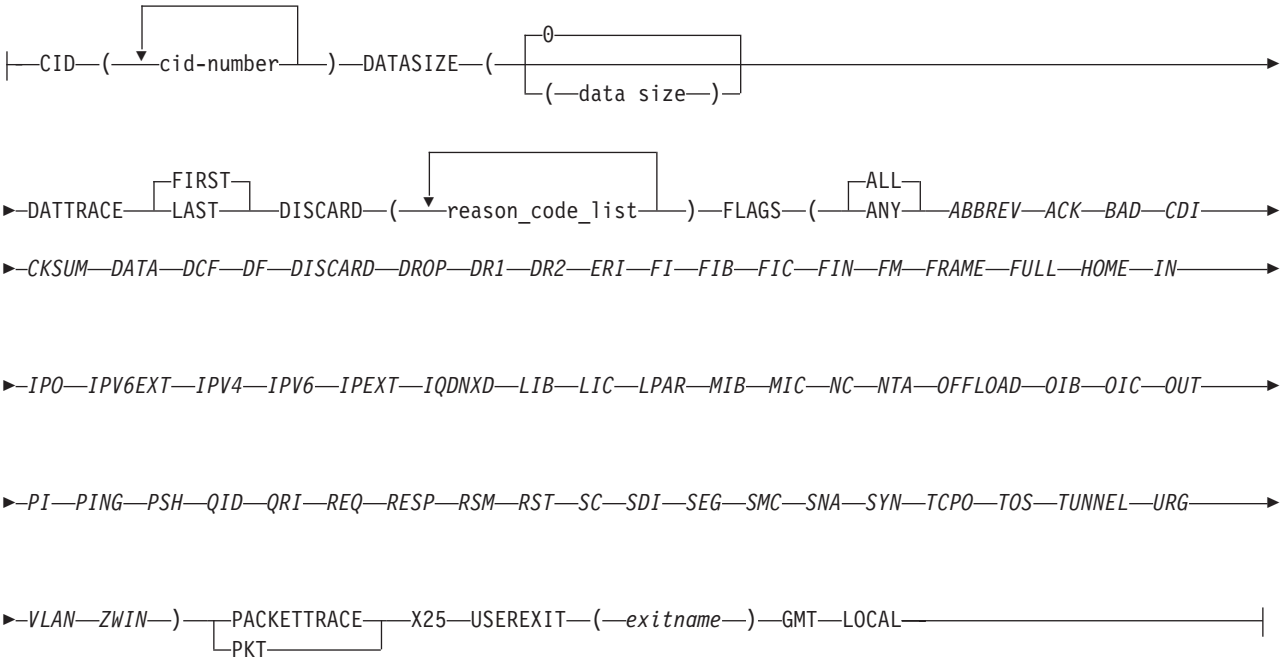

SNA Address:

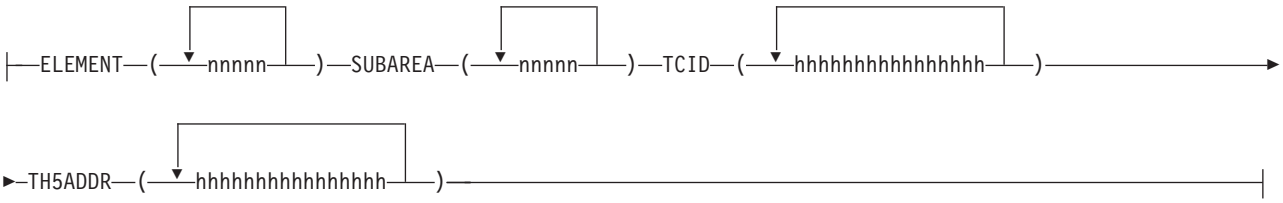

Report Generation:

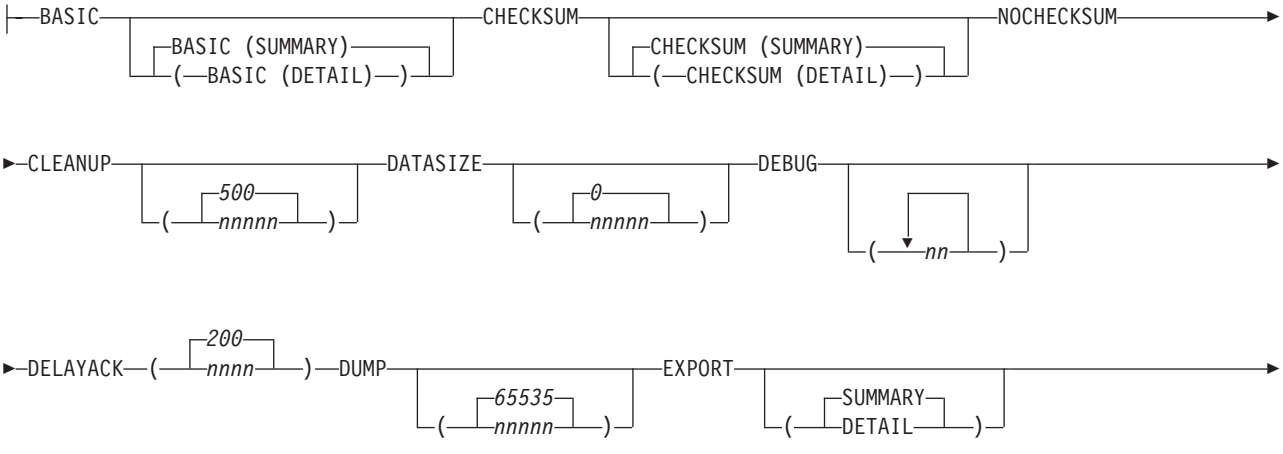

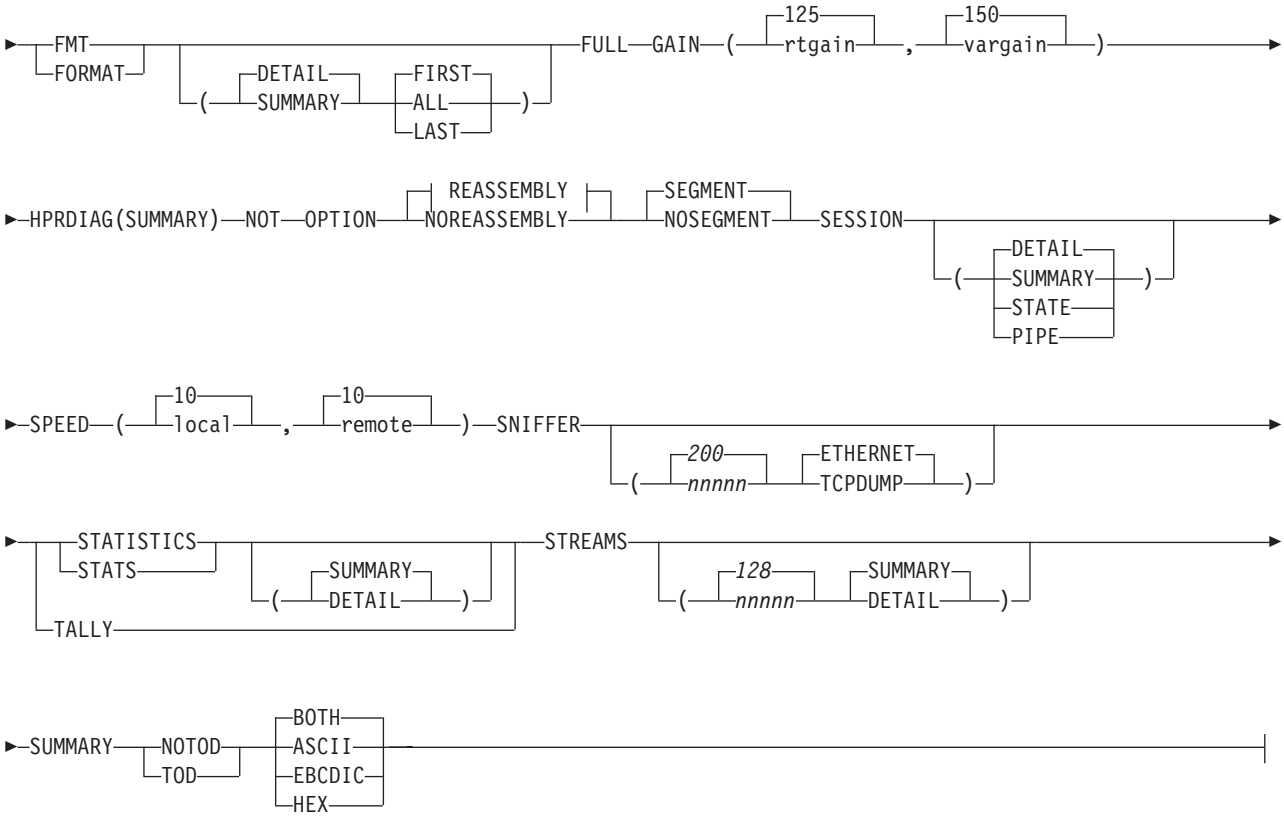

REASSEMBLY:

OPTIONS keywords

The following keywords are used for the OPTIONS component routine parameters. You can enter complete keywords. You can also enter a portion of the characters to make the keyword distinguishable from other keywords. For example, for the DISCARD keyword, you can enter DISCARD, DISCAR, DISCA, DISC, DIS, or DI.

AH Select packets with an AH extension header.

ASCII
Packet trace data dumped is shown in hexadecimal and interpreted in ASCII translation only. The default is BOTH.

BASIC ([DETAIL|SUMMARY])
For specific packet types, format each element of the packet data. This parameter applies to DNS, RIP, and SNMP packet data.

DETAIL
Format the IP header, protocol header, and protocol data in as few lines as possible. DETAIL is the default.

SUMMARY
Format the IP and protocol headers in as few lines as possible.

BOOTP[(port_number|67 port_number|68)]
Select BOOTP and DHCP protocol packets. The port_number defines the BOOTP and DHCP port numbers to select packets for formatting. Equivalent to PORT(67 68).

BOTH
Packet trace data dumped is shown in hexadecimal and interpreted with both ASCII and EBCDIC translations. The default is BOTH.

BROADCAST
Select packets with a broadcast IPv4 address. Equivalent to IPADDR(255.255.255.255/255.255.255.255).

CHECKSUM [(DETAIL|SUMMARY)]
The selected packets have their checksum values validated.

DETAIL
If there is a checksum error, then the packet is formatted and dumped.

SUMMARY
A message is issued for each packet that encounters a checksum error. SUMMARY is the default.

CID
Select data trace records that contain the specific connection ID value. The connection ID value can be determined from the Netstat COnn/-c report. For TCP connections across a Shared Memory Communications over Remote Direct Memory Access (SMC-R) link, data trace records can be selected by using the local SMC-R link ID (LocalSMCLinkId). This SMC-R link ID can be determined from the Netstat ALL/-A or Netstat DEvlinks/-d report. Up to 16 values or ranges can be specified.

CLASSA
Select packets with a class A IPv4 address. Equivalent to IPADDR(0.0.0.0/128.0.0.0).

CLASSB
Select packets with a class B IPv4 address. Equivalent to IPADDR(128.0.0.0/192.0.0.0).

CLASSC
Select packets with a class C IPv4 address. Equivalent to IPADDR(192.0.0.0/224.0.0.0).

CLASSD
Select packets with a class D IPv4 address. Equivalent to IPADDR(224.0.0.0/240.0.0.0).

CLASSE
Select packets with a class E IPv4 address. Equivalent to IPADDR(240.0.0.0/248.0.0.0).

CLEANUP(nnnnn|500)
Defines a record interval where saved packet information in storage is released. The minimum value is 500 records; the maximum value is 1048576 records; the default is 500 records. If you set the record interval to 0, cleanup does not occur.

DATASIZE (data_size|0)
Selects packets that contain more protocol data than the data_size value. The

minimum value is 0. The maximum value is 65535. The data size is determined from the amount of packet data available minus the size of any protocol headers. Equivalent to FLAGS(DATA).

DATTRACE
Select packets that are written from the VARY TCPIP,,DATTRACE command.

DEBUG(debug_level_list)
Provides documentation about SYSTCPDA format processing. debug_level_list is a list of numbers from 1 to 64. Use only under the direction of an IBM Service representative.

DELAYACK(threshold|200)
The delay acknowledgment threshold in milliseconds used in the calculation of round-trip time in the TCP session report. The minimum value is 10 milliseconds. The maximum value is 1000 milliseconds. The default value is 200 milliseconds.

DEVICEID(device_id)
Selects packets that are written to or received from an OSAENTA trace with one of the specified device identifiers. One to 16 device IDs can be specified. This filter applies only to type 7 trace records. The device_id value is a hexadecimal number in the form X'*csmfclua*':

cs The channel subsystem ID for this datapath device.

mf The LPAR Multiple Image Facility ID for the LPAR that is using this datapath device.

cl The control unit logical identifier for this datapath device.

ua The unit address for this datapath device.

Each identifier is a two-digit hexadecimal value in the range 00 - FF.

Tip: You can obtain the *device_id* values for any active user of the OSA by using the Hardware Management Console (HMC). For a data device that is active on a z/OS stack, you can obtain the *device_id* value for that data device from message IST2190I of the output from the D NET,TRLE command.

DEVTYPE(device_type_list)
Select packets that are written to or received from an interface with one of the specified device types. From 1 to 16 types can be specified. This does not apply to data trace records. The following types can be specified:

- ATM
- CDLC
- CLAW
- CTC
- ETHER8023
- ETHERNET
- ETHEROR8023
- FDDI
- HCH
- IBMTR
- IPAQENET
- IPAQENET6
- IPAQIDIO

- IPAQIDIO6
- IPAQTR
- IQDX
- IQDX6
- LOOPBACK
- LOOPBACK6
- MPCPTP
- MPCPTP6
- OSAFDDI
- OSAENET
- SNALINK
- SNALU62
- VIRTUAL
- VIRTUAL6
- X25NPSI

DISCARD(reason_code_list)
Select packets with one of the specified discard reason codes. Up to 16 discard reason codes can be specified in the range 0 - 65535. Each entry in the list can be a range: low_number:high_number. Values can be decimal or hexadecimal.

```
    0
      Packet was not discarded
    1:4087     A packet was discarded by OSA-Express
    1:1023     Select packets discarded by OSA-Express for DISCARD=EXCEPTION
reasons
    4096:8191  IP packet was discarded by TCPIP
    8192:12287 TCP packet was discarded by TCPIP
```

See z/OS Communications Server: IP and SNA Codes for the TCP/IP discard reason codes.

DNS[(port_number|53)]
Select Domain Namer Service protocol packets. The port_number defines the DNS port number to select packets for formatting. Equivalent to PORT(53).

DOMAIN[(port_number|53)]
Select Domain Namer Service protocol packets. The port_number defines the DNS port number to select packets for formatting. Equivalent to PORT(53).

DUMP[(nnnnn|65535)]
Dump the selected packets in hexadecimal with EBCDIC and ASCII translations. The IP and protocol headers are dumped separately from the packet data. The value *nnnnn* represents the maximum amount of packet data that is to be dumped from each packet. The default value is 65535 bytes. The minimum value is 0. The maximum value is 65535. The IP and protocol headers are not subject to this maximum.

The default report options are DUMP and FORMAT.

The BOTH, ASCII, EBCDIC, and HEX keywords describe how the dumped packets are translated. The default is BOTH. The display can be changed by using these keywords. The default ASCII translation table is used. This table might not match the table that is being used by the application. When you are formatting the CTRACE, it is helpful to have the correct line length. Use the IPCS PROFILE LINESIZE command to set the line length. For example,

```
IPCS PROFILE LINESIZE(80)
```

sets the maximum line length to 80 characters so that all formatted data is viewable within 80 characters.

If the STREAM report is chosen, then the dump of the packets is deferred until the stream of data is collected.

EBCDIC
Packet trace data dumped is shown in hexadecimal and interpreted with EBCDIC translation only. The default is BOTH.

EE Select Enterprise Extender (EE) protocol packets. The port number defines the first EE port number to select packets for formatting. The EE port number and the next four port numbers are used. Equivalent to PORT(12000:12004).

ELEMENT(element_number_list)
Select SNA protocol packets with a matching origin or destination element address in the TH2 or TH4 transmission header. Valid values are in the range 0 - 65535. Up to 16 element numbers can be specified.

ESP
Select packets with a protocol number of 50. Equivalent to PROTOCOL(50).

ETHTYPE(*type*)
Selects packets that are written to or received from an OSAENTA trace with one of the specified frame types. From 1 to 16 types can be specified. This filter applies only to type 7 trace records. The following types can be specified:

```
x'0800'  for IP
x'86DD' for IPV6
x'0806' for ARP
x'80d5' for SNA
```

EXPORT[(DETAIL|SUMMARY)]
The selected packets are written to the EXPORT data set in .CSV (Comma Separated Value) format. In .CSV format, each character field is surrounded by double quotation marks and successive fields are separated by commas. The first line of the file defines the fields. Each subsequent line is a record that contains the values for each field.

DETAIL
Format the IP header, protocol header, and protocol data as separate lines of data.

SUMMARY
Format the IP header and protocol header in one line of data. SUMMARY is the default.

Allocate a file with DDNAME of EXPORT before you invoke the CTRACE command with EXPORT in the OPTIONS string.

```
ALLOC
FILE(EXPORT) DA(PACKET.CSV) SPACE(15 15) TRACK
```

The record format is variable block with logical record length of 512 bytes.

FINGER[(port_number|79)]
Select FINGER protocol packets. The port_number defines the FINGER port number to select packets for formatting. Equivalent to PORT(79).

FIRST|LAST
Selects which packet in a set of encapsulated packets is used for selection. An example is the ICMP error report packet that contains the IP header that is in

error. FIRST indicates that the ICMP packet is used for selection. LAST indicates that the last encapsulated IP header is used for selection. FIRST is the default.

If a packet is encapsulated for IPSec with Encapsulating Security Payload (ESP), all inner packets are encrypted. In this case, FIRST is used for selection when these packets are analyzed.

FLAGS(flags list)
Select packets that have the matching characteristics. Flags that can be specified are:

ALL When more than one flag is specified, the packet must meet all the criteria of the flags requested. ALL is the default.

ANY When more than one flag is specified, the packet need meet only one of the criteria of the flags requested.

ABBREV
Select packets that are abbreviated.

ACK Select packets that have a TCP header with the ACK flag set.

BAD Select packets that might be too short to contain all the required headers

BBI The SNA packet contains a begin bracket indicator.

BCI The SNA packet contains a begin chain indicator.

CDI The SNA packet contains a change direction indicator.

CEBI The SNA packet contains a conditional end bracket indicator.

CKSUM
Select packets that have a check sum error

CLC The SMC Connection Layer Control packets

CSI The SNA packet contains a code selection indicator.

DATA Selects packets that contain data.

DF Select packets that have a non-zero discard code. These packets are discarded by TCP/IP.

DFC The SNA packet is a data flow control packet.

DISCARD
Select packets that have a non-zero discard code. These packets are discarded by OSA-Express or by TCP/IP.

DR1 The SNA packet is requesting a DR1 response.

DR2 The SNA packet is requesting a DR2 response.

EBI The SNA packet contains an end bracket indicator.

ECI The SNA packet contains an end chain indicator.

EDI The SNA packet contains an enciphered data indicator.

ERI The SNA packet is an error response.

FI The SNA packet contains formatted data.

FIB The SNA packet is the first packet of a bracket (or of a conditional begin bracket). The RH BBI flag is set and the EBI flag is not.

FIC Select packets that are the first in chain SNA RU.

FIN Select packets that have a TCP header with the FIN flag set.

FIS Select packets that are in the first fragment of an IPv4 or IPv6 packet or the first segment of a SNA PDU.

FMD The SNA packet is a function management data packet.

FMH The SNA packet is a function management data header.

FRAME
Selects OSAENTA packets that have a frame header.

FULL Select packets that are complete.

HOME
Select packets that have an IP destination address equal to the IP source address.

IN Select packets that are inbound.

IPEXT Select packets that have an extension header.

IPO Select packets that have an IPv4 header options field.

IPV4 Select IPv4 packets. IPv4 cannot be used in combination with other data selectors that are IPv6-specific, such as LINKLOCAL.

IPV6 Select IPv6 packets. IPv6 cannot be used in combination with other data selectors that are IPv4-specific, such as BROADCAST.

IPV6EXT
Select packets that have an extension header. IPV6EXT is equivalent to IPEXT.

IQDXND
Select ICMPV6 Neighbor Advertisement and Neighbor Solicit packets on an IQDX device.

LIB The SNA packet is a last packet of a bracket. The RH BBI flag is not set and the EBI flag is set.

LIC Select packets that are the last in a chain of SNA RUs.

LIS Select packets that are the last fragment of an IPv4 or IPv6 packet or the last segment of a SNA PDU.

LPAR Select NTA packets that are transmitted between LPARs shared by an OSA-Express device.

L2 The OSAENTA packet is from a layer 2 OSA application.

L3 The OSAENTA packet is from a layer 3 OSA application (like TCP/IP).

MIB The SNA packet is in the middle of a bracket. The RH BBI flag is not set and the EBI flag is not set.

MIC Select packets that are the middle fragment of an IPv4 or IPv6 packet.

MIS Select packets that are the middle fragment of an IPv4 or IPv6 packet or the middle segment of a SNA PDU.

NC The SNA packet is a Network Control packet.

NTA Select OSAENTA packets.

OFFLOAD
Select outbound packets for which segmentation is offloaded.

OIB The SNA packet is the only packet of a bracket. The RH BBI flag is set and the EBI flag is set.

OIC Select packets that are only in a chain SNA RH request.

OIS Select packets that are IPv4 or IPv6 packets that are not fragmented or that are the only segment of a SNA PDU.

OUT Select packets that are outbound.

PDI Select SNA packets with the padded data indicator.

PDU The IP packets that are packed by TCP/IP into a single PDU buffer.

PI The SNA packet contains a pacing indicator.

PING Select packets that are ICMP/ICMPv6 echo request and echo reply.

PSH Select packets that have a TCP header with the PSH flag set.

QID Select packets that have a QID value greater than one.

QRI The SNA packets with a queued response indicator

REQ The SNA packet is a request.

RESP The SNA packet is a response.

RLWS Select SNA packets with the request large window size indicator.

RSM Select packets that are reassembled.

RST Select packets that have a TCP header with the RST flag set.

SC The SNA packet is a session-control packet.

SDI The SNA packet contains sense data.

SEG Select packets that are segmented.

SMC Select SMC packets.

SNA Select SNA packets.

SYN Select packets that have a TCP header with the SYN flag set.

TCPO Select packets that have a TCP header options field.

TOS Select IPv4 packets that have a nonzero value in the ip_tos field.

TUNNEL
Select packets with protocol number 47 GRE or 41 (IPv6 over IPv4). z/OS Communications Server currently does not support IPv6 over IPv4 (protocol number 41).

URG Select packets that have a TCP header with the URG flag set.

VLAN Select packets that have a VLAN 802.1q tag

ZWIN Select packets that have a TCP header with a zero window value.

Notes:

- The use of the FIC, MIC, and LIC flags require the use of the NOREASSEMBLY option.
- When a packet is reassembled, then it becomes an OIC packet with the RSM flag set.
- Do not intermix SNA and IP flags.

Table 14. Flags that apply to IP or SNA packets

Flag	Applies to IP	Applies to SNA	Comments
ABBREV	Y	Y	
BAD	Y	Y	
BBI	N	Y	
BCI	N	Y	
CDI	N	Y	
CEBI	N	Y	
CI	N	Y	
CKSUM	Y	N	
CLC	Y	N	
CSI	N	Y	
DATA	Y	Y	
DF	Y	N	
DFC	N	Y	
DISCARD	Y	Y	
DR1	N	Y	
DR2	N	Y	
EBI	N	Y	
ECI	N	Y	
EDI	N	Y	
ERI	N	Y	
FI	N	Y	
FIB	N	Y	
FIC	N	Y	
FIN	Y	N	TCP only
FIS	Y	Y	
FM	N	Y	
FMD	N	Y	
FMH	N	Y	
FRAME	N	Y	OSAENTA only
FULL	Y	Y	
HOME	Y	N	
IN	Y	Y	
IPEXT	Y	N	
IPO	Y	N	
IPV4	Y	N	
IPV6	Y	N	
IPV6EXT	Y	N	
LIB	N	Y	
LIC	N	Y	
LIS	Y	Y	

Table 14. Flags that apply to IP or SNA packets (continued)

Flag	Applies to IP	Applies to SNA	Comments
LPAR	Y	Y	OSAENTA only
L2	Y	Y	OSAENTA only
L3	Y	Y	OSAENTA only
MIB	N	Y	
MIC	N	Y	
MIS	Y	Y	
NC	N	Y	
NTA	Y	Y	OSAENTA only
OFFLOAD	Y	N	TCP only
OIB	N	Y	
OIC	N	Y	
OIS	Y	Y	
OUT	Y	Y	
PDI	N	Y	
PDU	Y	N	SYSTCPDA only
PI	N	Y	
PING	Y	N	
PSH	Y	N	TCP only
QID	Y	Y	OSA-Express 3 or later ports with QDIO inbound workload queueing enabled.
QRI	N	Y	
REQ	N	Y	
RESP	N	Y	
RLWS	N	Y	
RSM	Y	N	
RST	Y	N	TCP only
SC	N	Y	
SDI	N	Y	
SEG	Y	Y	
SMC	Y	N	
SYN	Y	N	TCP only
TCPO	Y	N	TCP only
TOS	Y	Y	SNA TPF field
TUNNEL	Y	Y	
URG	Y	N	TCP only
VLAN	Y	Y	OSAENTA only
ZWIN	Y	N	TCP only

FMT
Equivalent to FORMAT.

FORMAT[(DETAIL|SUMMARY ALL|FIRST|LAST)]
The selected packets with defined packet data are to be formatted. The SHORT keyword on the CTRACE command selects this option if no other report options are specified. The default report options are DUMP and FORMAT.

DETAIL
Format the IP header, protocol header, and the protocol data.

SUMMARY
Format the IP header and protocol header. DETAIL is the default.

ALL
Format all encapsulated packets. ALL is the default.

FIRST
Format the first encapsulated packet.

LAST
Format the last encapsulated packet

An example of an encapsulated packet is an ICMP error report.

FTP[(data_port_number|20 control_port_number|21)]
Select FTP protocol packets. The port_number defines the FTP port numbers to select packets for formatting. Equivalent to PORT(20,21).

FULL
Equivalent to DUMP and FORMAT. The FULL keyword on the CTRACE command selects this option if no other report options are specified.

GAIN(rtgain|125,vargain|250)
Values of the round-trip gain (rtgain) and the variance gain (vargain), in milliseconds, used in the calculation of round-trip time in the TCP session report. Valid values are in the range 0 - 1000. The default value for rtgain is 125. The default value for vargain is 250.

GOPHER[(port_number|70)]
Select GOPHER protocol packets. The port_number defines the GOPHER port numbers to select packets for formatting. Equivalent to PORT(70).

GRE
Select packets with a protocol number of 47. Equivalent to PROTOCOL(47).

GMT
Format the time stamps in GMT time. The default is the value that is specified on the CTRACE subcommand.

HEX
Packet trace data dumped is shown in hexadecimal only with no translation. The default is BOTH.

HPRDIAG[(SUMMARY)]
Select high-performance routing (HPR) packets and group them by transport connection identifier (TCID). The report shows session information that can be helpful for HPR diagnosis.

HOST
Select packets with a host IP address. Equivalent to IPADDR(0.0.0.0/255.255.0.0)

HTTP[(port_number|80)]
Select HTTP protocol packets. The port_number defines the HTTP port numbers to select packets for formatting. Equivalent to PORT(80). See "WWW[(port_number|80)]" on page 120.

ICMP
Select packets with a protocol number of 1. Equivalent to PROTOCOL(1).

ICMP6 or ICMPV6
Select packets with a protocol number of 58. Equivalent to PROTOCOL(58).

IGMP
Select packets with a protocol number of 2. Equivalent to PROTOCOL(2).

INTERFACE(interface_name_list) or LINKNAME(interface_name_list)
Select packet trace records with the specified interface name. Up to 16 interface names can be specified. Each interface name can be up to 16 characters. Use an asterisk (*) as a wildcard to replace characters at the end of the interface name.

IPADDR(ipaddr[/mask_or_prefixlength]|X'hhhhhhhh'[]-nnnnn[)
Select packets with a matching IP address, optional IPv4 address mask or IPv6 prefix length and optional port number. Up to 16 IP addresses can be specified. The IPADDR is specified in three parts:

1. An IPv4 or IPv6 address

   The IPv4 address can be in dotted decimal notation, a keyword, or a hex value.

   - IPv4 dotted decimal notation

     `127.0.0.1`
   - IPv4 keyword

     A A class A IPv4 address, 0.0.0.0/128.0.0.0

     B A class B IPv4 address, 128.0.0.0/192.0.0.0

     C A class C IPv4 address, 192.0.0.0/224.0.0.0

     D A class D IPv4 address, 224.0.0.0/240.0.0.0

     E A class E IPv4 address, 240.0.0.0/248.0.0.0

     H A local host address, 0.0.0.0/0.0.255.255

     L An IPv4 or IPv6 loopback address, 127.0.0.0/255.0.0.0 or ::1

     M The broadcast IPv4 address, 255.255.255.255/255.255.255.255

     ***** Any address, 0.0.0.0/0.0.0.0

     0 An IPv4 or IPv6 address of zero, 0.0.0.0/255.255.255.255 or ::/128
   - IPv4 or IPv6 address as a hexadecimal number up to 32 (IPv4) or 128 (IPv6) digits

     `X'7f000001'`
   - IPv6 address

     `1080::8:800:200C:417A`
2. An IPv4 address mask or IPv6 prefix length

   The IPv4 address mask (1 - 32) or IPv6 prefix length (1 - 128) is preceded by a slash(/). Specify an IPv4 address mask only when the IPv4 address is in dotted decimal notation. The IPv4 address mask can be in dotted decimal notation, for example: `9.37/255.0.0.0` or `9.37/255.255.0.0`
3. A port number

The port number is preceded by a dash (-). It is a decimal number in the range 0 - 65535.

Notes:

- There should be no spaces between the IP addresses and the subnet masks.
- The BROADCAST, CLASSA, CLASSB, CLASSC, CLASSD, CLASSE, HOST, LINKLOCAL, LOOPBACK, MULTICAST, and SITELOCAL keywords add to the total of 16 IP addresses.
- The port number when used adds to the total of 16 port numbers in the PORT keyword.
- IPv4 addresses and IPv4 - mapped IPv6 addresses are treated as equivalent addresses.

IPID(ipid_number_list)
Select packets that match the ip_id number in the IPv4 packet header. Up to 16 ID numbers can be specified in the range 0 - 2147483647 or 0 - X'FFFFFFF'. Each entry in the list can be a range: low_number:high_number. Values can be decimal (nnnnn) or hexadecimal (X'hhhh'). If the packets are fragmented, specify NOREASSEMBLY to select each packet.

Tip: Associated encrypted text is not readable.

IPv4
Equivalent to FLAGS(IPV4).

IPv6
Equivalent to FLAGS(IPV6).

IKE
Select ISAKMP protocol packets. Equivalent to PORT(500). See the ISAKMP keyword.

ISAKMP
Select ISAKMP protocol packets. Equivalent to PORT(500). See the IKE keyword.

JOBLIST|JOBNAME(job_name_list)
Select data trace records with the specified JOBNAME. Up to 16 job names can be specified. Each job name can be up to eight characters. If the last character of a job name is an asterisk (*), then only the characters up to the asterisk are compared.

The CTRACE JOBLIST/JOBNAME parameter provides the same function, except that wildcards are not supported.

LIMIT(record_count)

record_count
The maximum number of records that are formatted. The default value 999999999 records.

Guideline: This keyword is also accepted if specified on the CTRACE subcommand.

LINKLOCAL
Select packets with an IPv6 link-local unicast prefix. Equivalent to IPADDR(FE80::/10).

LINKNAME(link_name_list)
Select packet trace records with the specified LINKNAME. Up to 16 link names

can be specified. Each link name can be up to 16 characters. If the last character of a link name is an asterisk (*), then only the characters up to the asterisk are compared.

The CTRACE JOBLIST/JOBNAME parameter provides the same function, except that wildcards are not supported and only the first eight characters of the link name are compared.

LOCAL
Format the time stamps in local time. The default is the value that is specified on the CTRACE subcommand.

LOOPBACK
Select packets with either an IPv4 or IPv6 loop back address. Equivalent to IPADDR(127.0.0.0/255.0.0.0::1). If other addresses are defined as loopback, they can be selected explicitly by using IPADDR().

LOOPBACK6
Select packets with an IPv6 loop back address. Equivalent to IPADDR(::1). If other addresses are defined as loopback, they can be selected explicitly by using IPADDR().

MACADDR(macaddr)
Selects packets that are written to or received from an OSAENTA trace with one of the specified MAC addresses. From 1 to 16 addresses can be specified. This filter applies only to type 7 trace records. A MACADDR is 12 hexadecimal digits.

MULTICAST
Select packets with either an IPv4 or IPv6 multicast address. Equivalent to CLASSD IPADDR(FF00::/8).

NAT
Select NAT protocol packets. Equivalent to PORT(4500).

NOCHECKSUM
The selected packets do not have their checksum values validated. CHECKSUM is the default.

NOREASSEMBLY
Do not reassemble fragmented IP packets into a complete packet. REASSEMBLY is the default.

NOSEGMENT
Packet trace records that span multiple CTRACE records are not recombined. Only the first segment record of packet is used. The rest of the segment records are discarded. SEGMENT is the default.

NOT
If the NOT option is selected then any selection criteria is reversed. If a record matches the selection criteria, it is not processed. If a record does not match the selection criteria, it is processed.

NTP[(port_number|123)]
Select NTP protocol packets. The port number defines the NTP port number to select packets for formatting. Equivalent to PORT(123).

OPTION
The selected options with defaults are listed.

OSPFI
Select packets with a protocol number of 89. Equivalent to PROTOCOL(89).

PACKETTRACE
Select packets that are written from the VARY TCPIP,,PKTTRACE command.

IPEXT
Select packets with an extension header.

PORT(port_number_list)
Select packets with one of the specified port numbers. Up to 16 port numbers can be specified in the range 0 - 65535. Each entry in the list can be a range: `low_number:high_number`. Values can be decimal (nnnnn) or hexadecimal (X'hhhh'). The following keywords add to the list of 16 port numbers:

- BOOTP
- DHCP
- DNS
- DOMAIN
- EE
- FINGER
- GOPHER
- HTTP
- NAT
- IKE
- RIP
- NTP
- ROUTER
- RPC
- SASP
- SMTP
- SNMP
- TELNET
- TFTP
- TIME
- WWW

PROTOCOL(protocol number list)
Select packets with one of the specified protocol numbers. Up to 16 protocol numbers can be specified in the range 0 - 255. Each entry in the list can be a range: `low_number:high_number`. Values can be decimal (nnn) or hexadecimal (X'hh').

Protocol filters on only the upper-layer header of an IPv6 packet. It does not filter for IPv6 extension headers (Hop-by-Hop Options, Routing, Fragment). Instead, IPv6 extension headers are included in the display of the basic IPv6 header. The following keywords add to the list of 16 protocol numbers:

- AH
- ESP
- GRE
- ICMP
- ICMP6,
- ICMPV6
- IGMP
- OSPFI

- TCP
- UDP

QOS(quality_of_service_list)
Select the records with the matching quality of service from the IPv4 Type of Service field. Up to 16 QoS values can be specified in the range 0 - 7. Each entry in the list can be a range: `low_number:high_number`. Values can be decimal (n) or hexadecimal (X'h').

QID(qid_list)
Select the records with the matching read queue identifier (QID) from the OSA-Express 3 or later ports with QDIO inbound workload queueing enabled. QID 1 selects records that are received on the primary input queue, and subsequent QIDs select records from the corresponding ancillary input queue (AIQ). Up to 16 QID values can be specified in the range 0 - 8. Each entry in the list can be a range: `low_number:high_number`. Values can be decimal (n) or hexadecimal (X'h').

REASSEMBLY[(packet_size|65535,DETAIL|SUMMARY)]
Reassemble IP fragments into a complete packet.

packet_size
The maximum size of a reassembled packet that is allowed. The smallest value that is allowed is 576 bytes, the largest is 65535 bytes. The default value is 65535 bytes.

DETAIL
List each of the reassembly statistics for each packet when a packet completes reassembly.

SUMMARY
Show only the reassembly statistics and information about packets that did not complete reassembly.

REASSEMBLY(65535,SUMMARY) is the default.

RECORDS(record_number_list)
Select the records with matching record numbers in the trace data. Up to 16 record numbers can be specified. Record numbers are assigned after any IPCS CTRACE selection criteria are met. Each entry in the list can be a range: `low_number:high_number`. Values can be decimal (nnnnnnnnnn) or hexadecimal (X'hhhhhhhh').

RIP[(port_number|520)]
Select RIP protocol packets. The port_number defines the RIP port number to select packets for formatting. Equivalent to PORT(520).

ROUTER[(port_number|520)]
Select RIP protocol packets. The port_number defines the RIP port number to select packets for formatting. Equivalent to PORT(520).

RIPNG
Select packets with a port number of PORT(521). Equivalent to PORT(521).

RPC[(port_number|111)]
Select RPC protocol packets. The port_number defines the RPC port number to select packets for formatting. Equivalent to PORT(111).

SASP (port_number|3860)
Select z/OS Load Balancing Advisor port numbers. The port_number defines the SASP port number to select packets for formatting. Equivalent to PORT(3860).

SEGMENT
Packet trace records that span multiple CTRACE records are recombined. Data from segment records is saved until all the CTRACE records are read to recreate the original packet. SEGMENT is the default.

SESSION[(DETAIL|PIPE|STATE|SUMMARY)]
Generate a report that shows TCP or UDP session traffic.

DETAIL
List each of the packets for a session, as well as the summary statistics. DETAIL is the default.

PIPE
List the amount of data left unacknowledged.

STATE
List the beginning and ending state of each session.

SUMMARY
Show only the summary statistics.

Tip: The UDP session analysis is also used for other protocols.

SITELOCAL
Select packets with an IPv6 site-local unicast address prefix. Equivalent to IPADDR(FEC0::/10).

SMCLLC
Select packets with a protocol number of 252. Equivalent to PROTOCOL(252).

SMTP[(port_number|25)]
Select SMTP protocol packets. The port_number defines the SMTP port number to select packets for formatting. Equivalent to PORT(25).

SNIFFER[(*nnnnn*|200, ETHERNET|TCPDUMP)]
Writes the trace records in a format acceptable for downloading to other trace analysis programs, such as programs from http://www.tcpdump.org/.

nnnnn
The maximum size of trace data. Packets with more data than this value are truncated. The default is 200 bytes. The largest value is derived from the LRECL of the SNIFFER data set.

ETHERNET
If this keyword is specified, the output is formatted for the Ethernet analysis application of the analyzer. This keyword specifies the file format only and does not imply that only packets traced on an Ethernet are collected. Packets from all devices can be collected by using this option.

The default for the SNIFFER option is ETHERNET.

TCPDUMP
The format is compatible with the files with an Ethernet header.

Note: The TOKENRING keyword on the CTRACE OPTIONS((SNIFFER(TOKENRING))) on the IPCS CTRACE subcommand is ignored. The ETHERNET format of the sniffer data set is selected.

The trace records are written to the file with a DD name of SNIFFER. After the file is generated, it can be downloaded as a binary file to the analyzer and loaded by using the standard features of the analyzer. Use NOREASSEMBLY to prevent the formatter from reassembling packets. Then, each packet is passed

as the packets are collected. The logical record length of the SNIFFER data set determines the largest amount of packet data that is written to the data set.

Allocate a file with DDNAME of SNIFFER before you invoke the CTRACE command with SNIFFER in the OPTIONS string as follows:

```
ALLOC FILE(SNIFFER)
DA(PACKET.TRC) SPACE(15 15) TRACK +
                      LRECL(8000) BLKSIZE(32000)
```

The data set has a record format of variable blocked with a logical record length of 8000 bytes. The maximum IP packet size is 7962 (8000 - 38) for SNIFFER(ETHERNET).

The minimum logical record length of the data set is 256 bytes.

Restriction: Do not use the SNIFFER option when the CTRACE subcommands are used with the IPCS MERGE subcommand. The SNIFFER data file is written over by the multiple CTRACE commands that specify the SNIFFER option.

SNMP[(port_number|161 port_number|162)]
Select SNMP protocol packets. The port_number defines the SNMP port number to select packets for formatting. Equivalent to PORT(161 162).

SPEED(local|10,remote|10)
The link speed, in megabits per second, for the local and remote link. These values are used in throughput calculations in the TCP session report. Valid values are in the range 0 - 17171. The default value is 10. Specify the slowest speed of the link in the route.

STATISTICS[(DETAIL|SUMMARY)]
After all the records are processed, generate statistical reports.

DETAIL
Reports are produced showing the number of records that are selected by record type, device type, job name, link name, protocol number, IP address and port numbers. The session summary report is a listing of the IP address and port number pairs that shows the number of records, the first and last record numbers, and the first and last record times.

SUMMARY
Only the session summary report is produced. SUMMARY is the default.

TALLY on the CTRACE command selects this option if no other report options are specified.

STATS
Equivalent to the STATISTICS option.

STREAMS[(stream_size|128 DETAIL|SUMMARY)]
Collect the packet data for dumping or formatting after the trace file is processed. The value *nnn* represents the maximum amount of storage that is used to capture each stream. The value *stream_size* represents the maximum amount of storage that is used to capture each stream. The smallest value is 16 KB. The largest value is 512 KB. The default value is 128 KB. The value is in 1024 bytes (1K) units.

SUMMARY
List about each packet in the stream. SUMMARY is the default.

DETAIL
Issue messages about the status of the stream.

Requirement: The DUMP keyword is required to dump the packet data.

SUBAREA(subarea_number_list)
Select SNA protocol packets with a matching subarea address in the TH4 transmission header. Valid values are in the range 1 - 65535. You can specify up to 16 subarea numbers.

SUMMARY
Format a single line for each trace record. SUMMARY on the CTRACE command selects this option if no other report options are specified. If no other report option specified on the CTRACE command, then SUMMARY is selected as the report.

NOTOD
Use the time that the trace data was moved to the CTRACE buffers for the reports. Normally, the time that the trace data was moved to the trace buffer is shown. The CTRACE command uses the time stamp when the trace data was moved to the buffers for START and STOP time selection. NOTOD is the default value for SYSTCPDA and SYSTCPIS traces.

TOD
Use the time that the trace data was captured for the reports. Normally, the time that the trace data was moved to the trace buffer is shown. The CTRACE command uses the time stamp when the trace data was moved to the buffers for START and STOP time selection. TOD is the default value for SYSTCPOT traces, which uses the time stamp that is generated by the OSA trace function.

TALLY
Equivalent to the STATISTICS(DETAIL) option.

TCID(transport connection_id_list)
Select SNA protocol packets with a matching transport connection identifier in the RTP transport header. Valid values include 1 - 16 hexadecimal digits. Up to 16 transport connection identifiers can be specified.

TCP
Select packets with a protocol number of 6. Equivalent to PROTOCOL(6).

TELNET[(port_number|23 [screen_width|80] [SUMMARY|DETAIL])]
Select TELNET protocol packets. The port_number defines the TELNET port number to select packets for formatting. Equivalent to PORT(23).

The screen_width parameter defines the value that is used for converting buffer offsets into row and column values for the 3270 data stream formatting. If the screen_width parameter is provided, then the port_number parameter must also be used. The minimum value is 80. The maximum value is 255. The default value is 80.

SUMMARY formats the 3270 data stream into a representation of the screen.

DETAIL formats each 3270 command and order.

There is no default for DETAIL or SUMMARY.

TFTP[(port_number|69)]
Select TFTP protocol packets. The port_number defines the TFTP port number to select packets for formatting. Equivalent to PORT(69).

TH5ADDR(session_address_list)
Select SNA protocol packets with a matching session address in the TH5 transmission header. Valid values include 1 - 16 hexadecimal digits. You can specify up to 16 session addresses.

TIME[(port_number|37)]
Select TIME protocol packets. The port_number defines the TIME port number to select packets for formatting.

TRAFFICCLASS(traffic_class)
Select the records with the matching IPv6 traffic class field. Up to 16 traffic class values can be specified in the range from 0 to 255. Each entry in the list can be a range: `low_number:high_number`. Values can be decimal (nn) or Hexadecimal (X'hh').

UDP
Select packets with a protocol number of 17. Equivalent to PROTOCOL(17).

USEREXIT(exitname)
Names the user exit to be called for each selected record. The USEREXIT keyword on the CTRACE command names a user exit that is called before the SYSTCPDA packet trace filtering is done. If this exit routine returns a nonzero return code, then the record is skipped by the SYSTCPDA formatter.

VLANID(*vlanid*)
Select packets that are written to or received from an OSAENTA trace with one of the specified VLAN identifiers. From 1 to 16 identifiers can be specified. This filter applies only to type 7 trace records. A VLAN identifier has a value in the range 0 - 4094.

Tip: The DEVICEID, MACADDR, ETHTYPE, and VLANID filter keywords apply to SYSTCPOT data. If these keywords are specified with SYSTCPDA data, then these filters are ignored.

WWW[(port_number|80)]
Select HTTP protocol packets. The port_number defines the HTTP port number to select packets for formatting. Equivalent to PORT(80).

X25
Select packet trace records created by the X25 processor.

Tip: This option is obsolete, but it is still accepted.

Report Examples

The CTRACE packet trace (SYSTCPDA) report generation outputs are described in the following examples.

Because IPv6 increases the IP address size, formatted IPv6 packet/data traces might be much wider than 80 columns.

OPTION:
Purpose

List the selected options and default keyword values.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((STAT(DETAIL) OPTION TCP))
```

```
COMPONENT TRACE SHORT FORMAT
 COMP(SYSTCPDA)SUBNAME((TCPCS))
 OPTIONS((STAT(DETAIL) OPTION TCP))
 z/OS TCP/IP Packet Trace Formatter, (C) IBM 2000-2005, 2004.365
 1  DSNAME('IPCS.R744334.DUMPA')

 2 OPTIONS((Both Bootp(67,68) Checksum(Summary) Cleanup(500) Datasize(0) DelayAck(200,200)
```

```
Domain(53) EE(12000:12004) Finger(79) Flags() Ftp(20,21) Gain(125,250) Gopher(70) Limit(999999999)
Gmt Ntp(123) Option Reassembly(65535,Summary) Router(520) Rpc(111) Sasp(3860) Segment Smtp(25)
Snmp(161,162) Speed(10,10) Statistics(Detail) Telnet(23,80,) Tftp(69) Time(37) Userexit() Www(80)
3  Protocol( /* 1 */ 6 /* TCP */, ) ))
```

The following fields are on the OPTION report.

1 DSNAME - The name of the source data.

2 OPTIONS((...)) - A listing of the active options with default values.

3 When a filter is specified, the list of filters with the number of filter values and filter values.

SUMMARY:
Purpose

Show one or two lines of information about each record in the trace.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SUMMARY
```

The following fields are on the SUMMARY report.

```
**** 2004/01/26
I - Inbound packet
O - Outbound packet

DP   Nr hh:mm:ss.mmmmmm IpId    Seq_num    Ack_num  Wndw Flags           DatLn Source/Destination
OT    1 14:18:00.447462 19E6 1452693653          0 32768 SYN                 0 10.7.1.61-3470
                                                                               192.168.248.44-5000
IT    2 14:18:00.601784 4E3C 3454024895 1452693654 32768 ACK SYN           0 192.168.248.44-5000
                                                                               10.7.1.61-3470
OT    3 14:18:00.601917 1A00 1452693654 3454024896 32768 ACK               0 10.7.1.61-3470
                                                                               192.168.248.44-5000
IT    4 14:18:01.111074 4E3D 3454024896 1452693654 32768 ACK PSH          47 192.168.248.44-5000
                        3232302D 2057656C *.......% 220- Wel*                  10.7.1.61-3470
IT    5 14:18:01.126148 4E3E 3454024943 1452693654 32768 ACK PSH          65 192.168.248.44-5000
                        32323020 52582E69 *........ 220 RX.i*                  10.7.1.61-3470
OT    6 14:18:01.126248 1A46 1452693654 3454025008 32703 ACK               0 10.7.1.61-3470
                                                                               192.168.248.44-5000
OT    7 14:18:03.290611 1B7F 1452693654 3454025008 32703 ACK PSH          10 10.7.1.61-3470
                        55534552 20637673 *........ USER cvs*                  192.168.248.44-5000
IT    8 14:18:03.373175 4E3F 3454025008 1452693664 32768 ACK PSH          32 192.168.248.44-5000
                        33333120 50617373 *....&/.. 331 Pass*                  10.7.1.61-3470
```

Figure 19. Example of a SUMMARY report

D Direction of the packet:

- **I** Inbound
- **O** Outbound

P The packet protocol

- **T** TCP
- **U** UDP
- **I** ICMP
- **G** IGMP
- **D** Data Trace

P Neither TCP, UDP, ICMP, nor IGMP

Nr The CTRACE record number

hh:mm:ss.mmmmmmm
The time stamp of the record

Source
The source IP address and port number

Destination
The destination IP address and port number

IpId
The packet ID number in hexadecimal

- For TCP

  seq_num
  The sequence number

  ack_num
  The acknowledgment sequence number

  wndw
  The window size

  flags
  The TCP header flags

  DatL
  The length of data in the datagram

  EBCDIC
  The first 8 bytes with EBCDIC translation

  ASCII
  The first 8 bytes with ASCII translation

- For UDP

  nnnnn
  The length of the UDP datagram

  DatL
  The length of the UDP packet data

  EBCDIC
  The first 8 bytes with EBCDIC translation

  ASCII
  The first 8 bytes with ASCII translation

- For ICMP

  ccccccccc
  The type of ICMP message

  xxxxxxxxxx
  The first 8 bytes of the user data in hexadecimal

- For IGMP

  nnnnn
  The maximum response time

  ccccccccc
  The type of IGMP message

nnn.nnn.nnn.nnn
The IGMP group address

- Other protocols

ccccccccc
The protocol name

nnnnn
The length of the protocol data

EXPORT:
Purpose

Reformat the information about the IP header, protocol header, and packet data into a file with CSV format.

Format

```
ALLOC FILE(EXPORT) DA(EXPORT.CSV) SPACE(15 15) TRACK

CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((EXPORT))
```

The following describe the EXPORT, EXPORT(SUMMARY), and EXPORT(DETAIL) report outputs.

- EXPORT

```
Export Report
 1  124 records written to USER2.EXPORT.CSV
 2  20,168 bytes written
```

The following fields are on the EXPORT report.

1
The number of data records written to the export data set.

2
The size of the export data set.

- EXPORT (SUMMARY)

```
"Flags ","Packet","Absolute Time  ","Rel Time","Delta Time",
"Device          ","Source          ","Destination     ",
"IpId","IpLen","Protocol  ","Summary"
"I O   ",       1,"19:49:42.788207",  0.000000,    0.000000,
"OSAQDIOLINK     ","9.67.115.17     ","9.67.115.63     ",
 17158,     78,"UDP","S=137 D=137 LEN=58"
"I O   ",      29,"19:52:21.240160",158.451952,    0.016739,
"OSAQDIOLINK     ","9.67.115.69     ","9.67.115.5      ",
  5971,     56,"ICMP","? LEN=28"
"I O   ",      37,"19:52:27.783944",164.995736,    0.000134,
"LOOPBACK        ","9.67.115.5      ","9.67.115.5      ",
   129,     56,"ICMP","? LEN=28"
"O O   ",      40,"19:52:39.284802",176.496595,    5.500260,
"OSAQDIO46       ","FEC9:C2D4::6:2900:EDC:217C","FEC9:C2D4::9:67:115:17",
    20,     60,"UDP","S=32810 D=33435 LEN=20"
"O O   ",      41,"19:52:39.284870",176.496662,    0.000067,
"OSAQDIO46       ","FEC9:C2D4::6:2900:EDC:217C","FF02::1:FF15:17",
    32,     72,"ICMPv6","ICMPv6"
"I O   ",      42,"19:52:39.285955",176.497748,    0.001085,
"OSAQDIO46       ","FEC9:C2D4::9:67:115:17","FEC9:C2D4::6:2900:EDC:217C",
    32,     72,"ICMPv6","ICMPv6"
"O O   ",      49,"19:52:58.286347",195.498140,   13.972912,
"LOOPBACK6       ","FEC9:C2D4::9:67:115:5","FEC9:C2D4::9:67:115:5",
    20,     60,"UDP","S=32810 D=33435 LEN=20"
"I O   ",      50,"19:52:58.286530",195.498323,    0.000182,
"LOOPBACK6       ","FEC9:C2D4::9:67:115:5","FEC9:C2D4::9:67:115:5",
    68,    108,"ICMPv6","ICMPv6"
```

The following describes fields found on the EXPORT (SUMMARY) report:

Control flags

Direction

- I — Input
- O — Output

A The packet was abbreviated (used with the following fragment flags).

R Reassembled packet.

O The Only fragment of a packet (it is complete).

F First fragment of a packet.

M Middle fragment of a packet.

L Last fragment of a packet.

T The packet was in a tunnel.

Packet
The packet number

Absolute Time
The time stamp on the packet

Rel Time
The time from the first packet in seconds

Delta Time
The time from the previous packet in seconds

Device
The device the packet was received on or sent from

Source
The source IP address

Destination
The destination IP address

IpId
The ID number from the IP packet header

IpLen
The length of the IP packet

Protocol
The protocol from the IP packet

Summary
Additional information from the protocol header.

- EXPORT (DETAIL)

```
"Flags ","Packet","Delta Time","Source         ",
"Destination     ","Protocol  ","Summary"
"I O   ",      10," 69.006010","9.67.115.5     ",
"9.67.115.5      ","IP"," S=9.67.115.5 D=9.67.115.5 LEN=71 ID=110"
"I O   ",      10," 69.006010","9.67.115.5     ",
"9.67.115.5      ","UDP"," S=1036 D=161 LEN=51"
"I O   ",      10," 69.006010","9.67.115.5     ",
"9.67.115.5      ","SNMP","GetRequest dpiPathNameForUnixStream.0"
"O O   ",      24,"  0.002956","9.67.115.5     ",
"9.67.115.69     ","IP"," S=9.67.115.5 D=9.67.115.69 LEN=40 ID=121"
"O O   ",      24,"  0.002956","9.67.115.5     ",
"9.67.115.69     ","UDP"," S=32810 D=33436 LEN=20"
```

```
"O O   ",      51,"  0.002695","FEC9:C2D4::9:67:115:5",
"FEC9:C2D4::9:67:115:5","IP"," S=FEC9:C2D4::9:67:115:5 D=FEC9:C2D4::9:67:115:5 LEN=60 ID=20"
"O O   ",      51,"  0.002695","FEC9:C2D4::9:67:115:5",
"FEC9:C2D4::9:67:115:5","UDP"," S=32810 D=33436 LEN=20"
"I O   ",      52,"  0.000244","FEC9:C2D4::9:67:115:5",
"FEC9:C2D4::9:67:115:5","IP"," S=FEC9:C2D4::9:67:115:5 D=FEC9:C2D4::9:67:115:5 LEN=108 ID=68"
"I O   ",      52,"  0.000244","FEC9:C2D4::9:67:115:5",
"FEC9:C2D4::9:67:115:5","ICMPv6"
```

The following describes fields found on the EXPORT (DETAIL) report:

Control flags

Direction

- I — Input
- O — Output

A The packet was abbreviated (used with the following fragment flags).

R Reassembled packet.

O Only fragment of a packet (it is complete).

F First fragment of a packet.

M Middle fragment of a packet.

L Last fragment of a packet.

T The packet was in a tunnel.

Packet
The packet number.

Delta Time
The time from the previous packet in seconds.

Source
The source IP address.

Destination
The destination IP address.

Protocol
There are multiple lines about a single packet. The first line contains "IP" to identify the data in the summary field. The second line identifies information about the protocol used by the packet. The possible third line identifies the application data in the packet.

Summary
Additional information from the protocol headers or packet data.

FORMAT:
Purpose

Format the CTRACE record header, the IP packet header, the protocol header, and the packet data. If one of the ports is a well-known port number and the SYSTCPDA supports data for the port number, the packet data is shown.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((FORMAT))
```

```
1   3 MVSJ     PACKET   00000001 23:39:11.873541 Packet Trace
2  To Interface      : TR1             Device: LCS Token Ring   Full=56
  Tod Clock        : 2002/02/12 23:39:11.873539
```

```
3  Sequence #       : 0               Flags: Pkt Ver2 Out
 Source Port      : 1025            Dest Port: 53    Asid: 001E TCB: 007F62C0
4  IpHeader: Version : 4              Header Length: 20
 Tos              : 00              QOS: Routine Normal Service
 Packet Length    : 56              ID Number: 000E
 Fragment         :                 Offset: 0
5   TTL              : 64           Protocol: UDP         CheckSum: A6FB FFFF
 Source           : 9.67.113.65
 Destination      : 9.37.80.3

6  UDP
 Source Port      : 1025  ()        Destination Port:  53    (domain)
 Datagram Length  : 36              CheckSum: AD0B FFFF
7  DNS: 28

================================================================================
8 ;; ->>DNS HEADER<<- opcode: QUERY,  status: NOERROR, id: 40266
;; flags: rd; Ques: 1, Ans: 0, Auth: 0, Addit: 0

;; QUESTIONS: 1
;; w3.ibm.com                          IN      AAAA
```

1

A summary line indicating the source of the trace record showing:

- The record number.
- The system name.
- The type of the trace record.
- The time the record was moved to the trace buffer, or with the TOD option the time the trace data was captured.
- The description of the trace record, Packet Trace, Data Trace, EE, or OSAENTA.

2

The trace header with these fields:

- The direction of the trace record: From or To.
- The network interface name (or job name for Data Trace).
- The device type.
- Full or Abbrev with amount of trace data available.
- The time the trace record was captured.
- The number of records lost.

3

The segment number of the CTRACE record for which this packet trace record was written.

The flags that describe this trace record:

CfTxt This packet contains confidential text.

Adj The CTRACE record size was adjusted to a halfword length.

Tunnel
This packet consists of multiple protocol headers, for example, UDP that has an SNA packet (EE).

Rsm This packet was reassembled from multiple packets.

Home The source and destination of this packet are the same.

IpExt The packet has IP header extensions.

LS The packet was part of a segmentation offload group of packets.

Ping The packet is ICMP or ICMPv6 echo or echo reply packet.

Dat The packet was written in response to a DATTRACE command.

AH The packet contains an authentication header.

Seg This trace record was assembled from multiple CTRACE record segments.

In This packet is an inbound packet.

Out This packet is an outbound packet.

Bad The packet had a validation error.

CkSum
The packet had a checksum error.

Pdus This packet was written when multiple packets were in the same set of buffers.

QID This packet contains a nonzero OSA QID.

Nta This packet was written in response to an OSAENTA command.

VLan This packet contains a VLAN identifier.

Frame This packet contains an Ethernet header.

LPAR This packet sent over an OSA as LPAR to LPAR.

L2 This packet was sent as a Level 2 packet (directly routed by using MAC addresses).

L3 This packet was sent as a Level 3 packet (indirectly routed by using IP addresses).

Dscrd This packet was discarded by OSA or by TCPIP.

Rsp This packet is a SNA RU response.

Req This packet is a SNA RU request.

FMD This packet is a SNA Function Management RU.

NC This packet is a SNA Network Control RU.

SC This packet is a SNA System Control RU.

DFC This packet is a SNA Data Flow Control RU.

FMH This packet contains a SNA Function Management header.

SMC This packet was sent over an SMC-R interface.

4

The IP header showing fields from the IPv4 header. The header length is the number of bytes for the header. The offset field is the number of bytes from the end of the IP header where the fragment appears. With the REASSEMBLY option active, this field always contains zeros.

5

The check sum value. If possible, the check sum of the packet is calculated. If the calculated value is X'FFFF', the check sum is correct. If X'0000', the check sum could not be calculated because the packet was incomplete or fragmented. Other values indicate a check sum error.

6

The UDP protocol header. The fields of the header are shown.

The port number for the PORT keywords define the port numbers to be used to invoke a format routine.

Port
Keyword

67, 68
BOOTP

67, 68
DHCP

53 Domain

12000:12004
EE

79 Finger

70 Gopher

80 HTTP

500
IKE

4500
NAT

123
NTP

111
RPC

520
RIP

520
Router

25 SMTP

161,162
SNMP

23 TELNET

69 TFTP

37 TIME

7

The length of the DNS packet data following is shown.

8

The DNS header and resource records are formatted. Using the protocol numbers and the known port numbers, format routines are invoked to format standard packet data records.

```
1   17 MVSN     PACKET   00000004 19:43:02.541728 Packet Trace
2  To Interface      : LOGETH5          Device: QDIO Ethernet    Full=6300
    Tod Clock        : 2004/10/18 19:43:02.541728                Intfx: 5
    Sequence #       : 0                Flags: Pkt Out Offl
3  IpHeader: Version : 4                Header Length: 20
    Tos              : 00               QOS: Routine Normal Service
    Offload Length   : 6300             ID Numbers: 0012-0016
    Fragment         :                  Offset: 0
4   TTL              : 64               Protocol: TCP            CheckSum: 0000 971D
    Source           : 8.1.1.1
    Destination      : 8.1.1.2

5  TCP
    Source Port      : 1026  ()         Destination Port: 1026  ()
    Sequence Number  : 3823117120       Ack Number: 3823533758
    Header Length    : 32               Flags: Ack Psh
    Window Size      : 32768            CheckSum: 120B 0000 Urgent Data Pointer: 0000
    Offload Segments : 4                Length: 1448             Last: 456
     Option          : NOP
     Option          : NOP
     Option          : Timestamp        Len: 10 Value: F3913448 Echo: F3913446
```

Figure 20. Format report example

1

A summary line indicating the source of the trace record showing:

- The record number.
- The system name.
- The type of the trace record.
- The time the record was moved to the trace buffer, or with the TOD option the time the trace data was captured.
- The description of the trace record, Packet Trace or Data Trace.

2

The trace header with these fields:

- The direction of the trace record: From or To.
- The network interface name (or job name for Data Trace).
- The device type.
- Full or Abbrev with amount of trace data available.
- The time the trace record was captured.
- The number of records lost.
- The packet trace header flags.
- For inbound traffic, the queue identifier (QID) and queue type (for example, SYSDIST, BULKDATA, or EE) on which this packet was received using QDIO inbound workload queueing. For more information about QIDs, see QDIO inbound workload queueing in z/OS Communications Server: IP Configuration Guide.

3

The IP header showing fields from the IPv4 header. The header length is the number of bytes for the header. The offset field is the number of bytes from the end of the IP header where the fragment appears. With the REASSEMBLY option active, this field always contains zeros. If segmentation is offloaded, the ID number field shows the range of IP identifiers represented by this send and the Offload Length field shows the total length of the send (total data length plus one set of headers).

4

The check sum value. If possible, the check sum of the packet is calculated. If the calculated value is X'FFFF', the check sum is correct. If X'0000', the check sum could not be calculated because the packet was incomplete or fragmented. Other values indicate a check sum error.

5

The TCP protocol header. The fields of the header are shown. If segmentation is offloaded, the Offload Segments field shows the number of TCP segments represented by this send and the length of each segment. The length of each segment is the data length (not including headers). If all the segments are the same size, then the Last field does not appear. If the remainder of data length is nonzero, then Last field contains the remainder.

DUMP:
Purpose

Format the IP header, protocol header, and packet data in hexadecimal. The data can also be translated into EBCDIC, ASCII, or both.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((DUMP))
```

```
1 MVS073   PACKET   00000001 19:49:42.788207 Packet Trace
 From Interface    : OSAQDIOLINK      Device: QDIO Ethernet    Full=78
  Tod Clock        : 2002/02/12 19:49:42.788204
  Sequence #       : 0                Flags: Pkt Ver2
  Source Port      : 137              Dest Port: 137   Asid: 002B TCB: 00000000

 1  IP Header          : 20
000000 4500004E 43060000 8011FEC2 09437311  0943733F

 2  Protocol Header    : 8
000000 00890089 003AD7D7

 3  Data               : 50     Data Length: 50
000000 AD3D0110 00010000 00000000 20464845 |.............. .=.......... FHE|
000010 50464345 4C454846 43455046 46464143 |&...<.....&..... PFCELEHFCEPFFFAC|
000020 41434143 41434143 41434142 4C000020 |............<... ACACACACACABL.. |
000030 0001                               |..              ..              |
```

1

The IP header is dumped with no translation.

2

The protocol header is dumped with no translation.

3

The packet data is dumped with the translation specified by the ASCII, BOTH, EBCDIC or HEX keyword. The default is BOTH. The amount of data dumped can be limited by the value specified with the DUMP keyword. The default is 65535 bytes.

HPR:
Purpose

Select HPR packets and group them by TCID.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) OPTIONS((HPRDIAG(SUMMARY)))
```

```
Hpr Stats Report
  16384 bytes of storage were used for the HPR reports

                                                PktOut     ByteOut        Pkt        Byte                        SlowDown
TCID                        Pkt        Byte     OfOrder    OfOrder       ReXmit     ReXmit       SlowDown         ElapsedTime
----                        ----       ------   -------    ---------    --------    ---------    --------     ---------------
1B7580F800010034               1         101          0           0           0           0           0     00:00:00.000000
1B7580F900010035            3789      220212          0           0           0           0           0     00:00:00.000000
1B762A7F0001001E               1         156          0           0           0           0           0     00:00:00.000000
1B762A800001001F            3340      218802          0           0           0           0           0     00:00:00.000000
1B762A800001002D            4340      518802          0           0           3         258           0     00:00:00.000000
```

TCID Provides the name of the TCID (local or remote) for that particular session.

Pkt Total number of packets for that TCID.

Byte Total number of bytes for that TCID.

PktOutOfOrder
Number of out-of-order packets for that TCID.

ByteOutOfOrder
Number of out-of-order bytes for that TCID.

PktReXmit
Number of packets retransmitted for that TCID.

ByteReXmit
Number of bytes retransmitted for that TCID.

SlowDwn
Number of ARB slowdowns for that TCID.

SlowDownElapsedTime
Total time spent in adaptive rate base (ARB) slowdown mode.

REASSEMBLY:
Purpose

This report shows the packets that were reassembled. Use the REASSEMBLY(DETAIL) option to see all the packets that were reassembled. If the reassembled packets are larger than 32K then use `REASSEMBLY(`*nnnnn*`)`, where *nnnnn* is the maximum size of a reassembled packet.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((REASSEMBLY(DETAIL) STAT))

1  Reassembly of: 9.67.113.65-0 9.27.11.173-0 Id: 0043 status: +Fis +Lis
2   Rcd Nr Time            Delta           Offset Length   Next   Gap  Data Flags
   1638 15:28:49.975479 00:00:00.000000      0   3976   3976     0  3976 Fis
   1639 15:28:49.975501 00:00:00.000022   3976   3976   7952     0  3976 Mis
   1640 15:28:49.975524 00:00:00.000044   7952   3976  11928     0  3976 Mis
   1641 15:28:49.975545 00:00:00.000066  11928   3976  15904     0  3976 Mis
   1642 15:28:49.975567 00:00:00.000088  15904   3976  19880     0  3976 Mis
   1643 15:28:49.975594 00:00:00.000115  19880   3976  23856     0  3976 Mis
   1644 15:28:49.975620 00:00:00.000141  23856   3976  27832     0  3976 Mis
   1645 15:28:49.975689 00:00:00.000210  27832   3976  31808     0  3976 Mis
   1646 15:28:49.975737 00:00:00.000258  31808   3976  35784     0  3976 Mis
   1647 15:28:49.975771 00:00:00.000292  35784   3976  39760     0  3976 Mis
   1648 15:28:49.975804 00:00:00.000325  39760   3976  43736     0  3976 Mis
   1649 15:28:49.975835 00:00:00.000356  43736   3976  47712     0  3976 Mis
   1650 15:28:49.975865 00:00:00.000386  47712   3976  51688     0  3976 Mis
   1651 15:28:49.975898 00:00:00.000419  51688   3976  55664     0  3976 Mis
   1652 15:28:49.975926 00:00:00.000447  55664   3976  59640     0  3976 Mis
   1653 15:28:49.975962 00:00:00.000483  59640   3976  63616     0  3976 Mis
   1654 15:28:49.975986 00:00:00.000507  63616    392  64008     0   392 Lis
 64,008 bytes is the final length of the IP packet
```

```
     17 packets were used for reassembly
  64,008 bytes were accumulated for reassembly
================================================================================

 3  Packet Reassembly Report
Maximum reassembly buffer size is 65535
 4  Reassembly of: 9.27.11.173-0 9.67.113.65-0 Id: 3694 status: +Fis +Lis
  Rcd Nr Time            Delta           Offset Length   Next   Gap  Data Flags
    1655 15:28:50.024685 00:00:00.000000      0   1480   1480     0  1480 Fis
    1656 15:28:50.024705 00:00:00.000019   1480   1480   2960     0  1480 Mis
    1657 15:28:50.024739 00:00:00.000053   2960   1480   4440     0  1480 Mis
    1658 15:28:50.024772 00:00:00.000086   4440   1480   5920     0  1480 Mis
    1659 15:28:50.025506 00:00:00.000820   5920   1480   7400     0  1480 Mis
    1660 15:28:50.030534 00:00:00.005848   7400   1480   8880     0  1480 Mis
    1661 15:28:50.030592 00:00:00.005906   8880   1480  10360     0  1480 Mis
    1662 15:28:50.030607 00:00:00.005921  10360   1480  11840     0  1480 Mis
    1663 15:28:50.030650 00:00:00.005964  11840   1480  13320     0  1480 Mis
    1664 15:28:50.030683 00:00:00.005997  13320   1480  14800     0  1480 Mis
    1665 15:28:50.030698 00:00:00.006012  14800   1480  16280     0  1480 Mis
    1666 15:28:50.042927 00:00:00.018241  16280   1480  17760     0  1480 Mis
    1667 15:28:50.042946 00:00:00.018261  17760   1480  19240     0  1480 Mis
    1668 15:28:50.043006 00:00:00.018320  19240   1480  20720     0  1480 Mis
    1669 15:28:50.043021 00:00:00.018335  20720   1480  22200     0  1480 Mis
    1670 15:28:50.043058 00:00:00.018372  22200   1480  23680     0  1480 Mis
    1671 15:28:50.043114 00:00:00.018428  23680   1480  25160     0  1480 Mis
    1672 15:28:50.043130 00:00:00.018444  25160   1480  26640     0  1480 Mis
    1673 15:28:50.043174 00:00:00.018488  26640   1480  28120     0  1480 Mis
    1674 15:28:50.043222 00:00:00.018536  28120   1480  29600     0  1480 Mis
    1675 15:28:50.043257 00:00:00.018571  29600   1480  31080     0  1480 Mis
    1676 15:28:50.043544 00:00:00.018858  31080   1480  32560     0  1480 Mis
    1677 15:28:50.043592 00:00:00.018906  32560   1480  34040     0  1480 Mis
    1678 15:28:50.043607 00:00:00.018921  34040   1480  35520     0  1480 Mis
    1679 15:28:50.044618 00:00:00.019932  35520   1480  37000     0  1480 Mis
    1680 15:28:50.044649 00:00:00.019963  37000   1480  38480     0  1480 Mis
    1681 15:28:50.044698 00:00:00.020012  38480   1480  39960     0  1480 Mis
    1682 15:28:50.044712 00:00:00.020026  39960   1480  41440     0  1480 Mis
    1683 15:28:50.044745 00:00:00.020059  41440   1480  42920     0  1480 Mis
    1684 15:28:50.044778 00:00:00.020092  42920   1480  44400     0  1480 Mis
    1685 15:28:50.050178 00:00:00.025492  44400   1480  45880     0  1480 Mis
    1686 15:28:50.050212 00:00:00.025527  45880   1480  47360     0  1480 Mis
    1687 15:28:50.050244 00:00:00.025558  48840   1480  50320 -1480  1480 Mis
    1688 15:28:50.050275 00:00:00.025589  50320   1480  51800     0  1480 Mis
    1689 15:28:50.050328 00:00:00.025642  51800   1480  53280     0  1480 Mis
    1690 15:28:50.050343 00:00:00.025657  53280   1480  54760     0  1480 Mis
    1691 15:28:50.054558 00:00:00.029872  54760   1480  56240     0  1480 Mis
    1692 15:28:50.054614 00:00:00.029928  57720   1480  59200 -1480  1480 Mis
    1693 15:28:50.054628 00:00:00.029942  59200   1480  60680     0  1480 Mis
    1694 15:28:50.054680 00:00:00.029994  62160   1480  63640 -1480  1480 Mis
    1695 15:28:50.054694 00:00:00.030008  63640    368  64008     0   368 Lis
  64,008 bytes is the final length of the IP packet
      41 packets were used for reassembly
  59,568 bytes were accumulated for reassembly
--------------------------------------------------------------------------------

 5  1,641 packets required reassembly
    54 IP packet reassemblies were done
    52 IP packets were completely reassembled
     2 IP packets were incomplete
     0 packets failed reassembly
 1,627 storage requests for buffers were made
64,080 bytes of buffer space are still in use
191,872 bytes of buffer space was the maximum in use
114,688 bytes of control storage were used
```

Reassembly is always done (except with the NOREASSEMBLY option). However, the REASSEMBLY(DETAIL) option is needed for the report on completed reassemblies.

1

The current packet that was reassembled is identified with source and destination IP address and port numbers. The IP identification number is shown. The status of the reassembly is shown. Completed packets are shown when the final packet is received. Incomplete packets are shown during the final processing.

2

Each packet that was reassembled is shown. The flag shows the type of packet:

Fis First in the segment. The offset was 0.

Mis Middle in the segment. The offset was a nonzero value and the more fragment flag was set.

Lis Last in the segment. The offset was nonzero and the more fragment flag was not set.

Ooo The packet arrived out of order.

The Gap field is the number of bytes between the end of one packet and the start of the next. This should have a value of zero for normal processing. Nonzero values indicate duplicate data being sent.

3

When all the trace records have been processed the final report on reassembly is formatted. The maximum reassembly buffer size is shown. Packets that would exceed this size are rejected. This simulates the Ping of Death processing.

4

Incomplete packets that did not complete reassembly are shown.

5

The total number of trace records that were reassembled is shown with other statistics.

200 packets required reassembly
The number of packets that required reassembly (that had a fragment offset or the more fragment flag set).

57 IP packet reassemblies were done
The number of reassembled packets.

54 IP packets were completely reassembled
The number of reassembled packets where all the fragments were found.

3 IP packets were incomplete
The number of reassembled packets where all the fragments were not found.

0 packets failed reassembly
The number of packets that would have caused the completed packet to exceed the reassembly size.

170 storage requests for buffers were made
The number of times a request for reassembly buffer was made.

128,747 bytes of buffer space is still in use
The amount of storage still in use for incomplete packets.

284,158 bytes of buffer space was the maximum in use
The maximum amount of storage in use while reassembling packets.

Guideline: For reassembled packets, the calculated check sum fields are not X'FFFF', because the packets were modified by the reassembly process.

SESSION TCP:
Purpose

This report shows traffic for a TCP session.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((SESSION TCP))
```

```
-------------------------------------------------------------------------------
1  2 packets summarized
 Local Ip Address:          FEC9:C2D4::6:2900:EDC:217C
 Remote Ip Address:         FEC9:C2D4::9:67:115:17
Host:                              Local,          Remote
 Client or Server:                CLIENT,          SERVER
 Port:                              1027,              21
 Application:                          ,             ftp
 Link speed (parm):                   10,              10 Megabits/s
2  Connection:
 First timestamp:                          19:55:46.934032
 Last timestamp:                           19:55:46.934989
 Duration:                                 00:00:00.000957
 Average Round-Trip-Time:                             0.000 sec
 Final Round-Trip-Time:                               0.000 sec
 Final state:                       CLOSED (PASSIVE RESET)
 Out-of-order timestamps:                               0
3  Data Quantity & Throughput:       Inbound,        Outbound
 Application data bytes:              0,               0
 Sequence number delta:               0,               1
 Total bytes Sent:                    0,               0
 Bytes retransmitted:                 0,               0
 Throughput:                          0,               0 Kilobytes/s
 Bandwidth utilization:            0.00%,           0.00%
 Delay ACK Threshold:               200,             200 ms
 Minimum Ack Time:              0.000957,        0.000000
 Average Ack Time:              0.000957,        0.000000
 Maximum Ack Time:              0.000957,        0.000000
4  Data Segment Stats:            Inbound,        Outbound
 Number of data segments:             0,               0
 Maximum segment size:             1432,               0
 Largest segment size:                0,               0
 Average segment size:                0,               0
 Smallest segment size:               0,               0
 Segments/window:                   0.0,             0.0
 Average bytes/window:                0,               0
 Most bytes/window:                   0,               0
 Offload Sends:                       3            ( 50%)
 Offload Segments:                    6
 Offload Bytes:                   43616           (72.69%)
 Total Packets(normal + offload):    18           (83.33%)
5  Window Stats:                   Inbound,        Outbound
 Number of windows:                   0,               0
 Maximum window size:                 0,               0
 Largest window advertised:           0,           32768
 Average window advertised:           0,           32768
 Smallest window advertised:          0,           32768
 Window scale factor:                 0,               0
 Window frequency:                    0,               0 Windows/s
 Time Stamp updates:                  0,               0
 Total Round Trip Time:         0.000000,        0.000000 (    0%), (    0%)
 Average Round Trip Time:       0.000000,        0.000000
```

```
6  Number of:                          Inbound,        Outbound
 Packets:                              1,              1
  (x) Untraced Packets:                0,              0
  (.) In-order data:                   0,              0 ( 0.00%), ( 0.00%)
  (a) Acknowledgments:                 1,              0 (100.00%), ( 0.00%)
  (+) Data and ACK:                    0,              0 ( 0.00%), ( 0.00%)
  (u) Duplicate ACKs:                  0,              0 ( 0.00%), ( 0.00%)
  (w) Window size updates:             0,              0 ( 0.00%), ( 0.00%)
  (z) Zero window sizes:               0,              0 ( 0.00%), ( 0.00%)
  (p) Window probes:                   0,              0 ( 0.00%), ( 0.00%)
  (k) Keepalive segments:              0,              0 ( 0.00%), ( 0.00%)
  (r) Retransmissions:                 0,              0 ( 0.00%), ( 0.00%)
  (o) Out-of-order:                    0,              0 ( 0.00%), ( 0.00%)
  (d) Delayed ACKs:                    0,              0 ( 0.00%), ( 0.00%)
  (f) Fragments:                       0,              0 ( 0.00%), ( 0.00%)
7  Time Spent on:                      Inbound,        Outbound
  (.) In-order data:        00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (a) Acknowledgments:      00:00:00.000957, 00:00:00.000000 (106.33%), ( 0.00%)
  (+) Data and ACK:         00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (u) Duplicate ACKs:       00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (w) Window size updates:  00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (z) Zero window sizes:    00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (p) Window probes:        00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (k) Keepalive segments:   00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (r) Retransmissions:      00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (o) Out-of-order:         00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (d) Delayed ACKs:         00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (f) Fragments:            00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
8  Number of:                          Inbound,        Outbound
  (   S ) SYN:                         0,              1 ( 0.00%), (100.00%)
  ( A S ) ACK SYN:                     0,              0 ( 0.00%), ( 0.00%)
  (    F) FIN:                         0,              0 ( 0.00%), ( 0.00%)
  ( A  F) ACK FIN:                     0,              0 ( 0.00%), ( 0.00%)
  (  R  ) RST:                         1,              0 (100.00%), ( 0.00%)
  (U    ) URG:                         0,              0 ( 0.00%), ( 0.00%)
9  Time Spent on:                      Inbound,        Outbound
  (    S ) SYN:             00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  ( A  S ) ACK SYN:         00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (     F) FIN:             00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  ( A   F) ACK FIN:         00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)
  (   R  ) RST:             00:00:00.000957, 00:00:00.000000 (106.33%), ( 0.00%)
  (U     ) URG:             00:00:00.000000, 00:00:00.000000 ( 0.00%), ( 0.00%)

Messages:
2) The largest inbound window is less than twice the inbound MSS.
2) This may reduce inbound throughput for bulk data transfers.
2) It is usually desirable for the window size to be at twice the MSS.
3) The outbound side of the connection appears to be a bulk data transfer.
        I - Inbound packet
        O - Outbound packet
10  TcpHdr  IO F        Seq        Ack RcvWnd  Data Delta Time      TimeStamp  RcdNr    State     Inf Ip_id   Rtt    TimeStmpV  Time
     S   O     29260429          0  32768     0   0.000000 19:55:46.934032    101    SYN_SENT   4  0028    0.00      0.000 316250
   A R   I a          0   29260430      0     0   0.000957 19:55:46.934989    102      CLOSED   4  0014    0.00      0.000 316250
```

1

Host

The number of packets records for this session; the IP addresses and port of the session.

2

Connection

The first and last time of the session, the length of the session, the final value of RRT, and the final state of the session.

3

Data Quality & Throughput

These statistics are about the quantity of data transmitted. The number of bytes received inbound and the number of bytes send outbound.

4

Data Segment Stats

These statistics are about the segments, the number of segments, and the sizes of the segments. The maximum segment size is captured from the SYN packet. Offload statistics appear only when there were any offloaded packets. These values reflect the number of offload packets, the number segments in these offloaded packets, the number of bytes in offloaded packets, and the total number of segments that are sent from the interface.

5

Window Stats

These statistics are about the window changes. The Window scale factor is captured from the SYN packet. The Time Stamp updates are captured from the Tcp header options.

6

Number of Packets

These statistics are about the number of data packets that flow for carrying data. The percentages are based on the number of packets.

7

Time Spent on:

These statistics are about the delta times of data packets that flow for carrying data. The percentages are based on the duration of the session.

8

Number of

These statistics are about the number of control packets that flow for starting and ending a session. The percentages are based on the number of packets.

9

Time Spent on

These statistics are about the delta times of control packets that flow for starting and ending a session. The percentages are based on the duration of the session.

10

Details TcpHdr

The flags from the TCP header

***** This packet is reassembled.

A This packet is an acknowledgment.

P This packet has the PUSH flag set.

U This packet is urgent.

S This packet is a syn.

F This packet is a fin.

R This packet is a reset.

The type of data packet has one of the following flags:

. The packet flowed in order as per its sequence number.

x There is a gap in the sequence number and it looks like there is untraced data.

a The packet is a stand-alone acknowledgment of previously received data.

+ The packet is an acknowledgment of previously received data and also contains data.

u The packet is an acknowledgment of data previously acknowledged.

w The packet updated the window size.

z The packet changed the window size to zero.

p The packet was a window probe.

k The packet was a keepalive packet.

r The packet was retransmission.

o The packet arrived out of order.

d The packet exceeded the delay time threshold.

f The packet was a fragment of a complete IP packet.

! A dropped packet that had a checksum error, that was a fragment, or that was discarded.

SESSION UDP:
Purpose

This report shows traffic for a UDP session.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((SESSION UDP))
```

```
-------------------------------------------------------------------------------
   1  2 UDP packets summarized
   2  Local  IP address:                9.42.105.112
    Local  port, app:                        1313
    Remote IP address:                9.42.106.4
    Remote port, app:                          53 domain
    Protocol:                                 UDP
   3  First TimeStamp:     2011/01/21 19:37:01.024385
    Last TimeStamp:      2011/01/21 19:37:01.027320
    Duration:                    00:00:00.002935
   4  Statistics:                Inbound,   Outbound
      Packets:                       1,          1
      Fragments:                     0,          0 (    0%), (    0%)
      Bytes sent:                  261,         28
      Time spent sending:    00:00:00.0, 00:00:00.0 (  100%), (    0%)
      Time spent fragments:  00:00:00.0, 00:00:00.0 (    0%), (    0%)
      Throughput:                90.512,       9.71 K/s
      Largest segment size:        261,         28
      Average segment size:        261,         28
      Smallest segment size:       261,         28
      Direction change:              1,          0
   5  UDP Flags IO           Delta            Time DatLn  RcdNr Inf Ip_id Froff Drop     Esp#
             O  00:00:00.000000 19:37:01.024385    28      7   9  0717     0    0        0
             I  00:00:00.002935 19:37:01.027320   261      8   9  68E4     0    0        0
```

1
The number of packet records for this session.

2

The IP addresses and port numbers of the session, and the protocol that is used for the session.

3

The first and last time of a packet in the session and the length of the session.

4

Data Segment Stats

These statistics are about the packets, the number of packets, and the sizes of the packet. The largest, average, and minimum packet size. The time spent sending and receiving packets.

5

Details about each packet. The following flags can be set:

D `Don't Fragment` flag is set

O Only in chain packet

F First in chain packet

M Middle in chain packet

L Last in chain packet

T Encapsulated packet

R Reassembled packet

P ICMP Echo packet

Q ICMP Echo request packet

R ICMP Echo reply packet

C Check sum error in packet

N OSAENTA packet

V VLAN ID in packet

2 Layer 2 packet

3 Layer 3 packet

! Discarded packet

I Inbound packet

O Outbound packet

SNIFFER:
Purpose

This report shows information that is written to the SNIFFER data set.

Restriction: Do not use the SNIFFER option when the CTRACE subcommands are used with the IPCS MERGE subcommand. The SNIFFER data file is written over by the multiple CTRACE commands that specify the SNIFFER option.

Format

```
ALLOC F(SNIFFER) DATASET(SNIFFER.TRC) LRECL(1600) RECFM(V B) REUSE TRACK SPACE(15 15)

CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((SNIFFER NOREASSEMBLY STATS))
```

```
Interface Table Report
 1  Index Count Link             Address
    1     5 OSAQDIOLINK      9.67.115.63
    2    42 LOOPBACK         127.0.0.1
    3    18 OSAQDIOLINK      9.67.115.5
    4    31 OSAQDIO46        FEC9:C2D4::6:2900:EDC:217C
    5    21 OSAQDIO46        FE80::6:2900:EDC:217C
    6     6 LOOPBACK6        ::1
================================================================================

Sniffer Report
 2  13,963 records written to USER2.SNIFFER.ETH
 3  1,730,182 bytes written
 4  1184 packets were abbreviated
 5  200 is the maximum data size
 6  12438 packets were truncated from 1546 bytes
```

1

The list of device names found in the selected records. Each device is assigned an interface index.

2

This record count includes the two header records and one trailer record written to the SNIFFER data set.

3

The number of data bytes written to the SNIFFER data set. This is the amount of data to be downloaded.

4

The number of abbreviated records. This number is included in **6**.

5

Maximum size of the truncated records.

6

The number of truncated records. Records were truncated because the size of the packet exceeded the logical record length of the SNIFFER file. Increase the logical record length to prevent the records from being truncated. The maximum logical record length is 32,763 or the size of physical disk blocks, whichever is smaller.

STATISTICS:
Purpose

The records are counted by record type, device type, device name, job name, protocol, IP address, TCP port number, and UDP port number.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((STATISTICS(DETAIL)))
```

```
 1    No packets required reassembly
================================================================================
SYSTCPDA Trace Statistics
   123 ctrace records processed
     0 segmented trace records read
     0 segmented trace records were lost
   123 trace records read
     0 records could not be validated
   123 records passed filtering
   123 packet trace records processed
     0 data trace records processed
================================================================================
 2  Record Type Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Record Type
     65     55       3814     10        644     1 2002/02/12 19:49:42   123 2002/02/12 19:57:45 1(Packet Trace)
     58     19       1828     39       2840    40 2002/02/12 19:52:39   117 2002/02/12 19:56:29 4(Packet Trace)
```

```
    123     74       5642     49       3484 Total
     2 Record Type(s) found
Ip Version Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Ip Version
     65     55       3814     10        644     1 2002/02/12 19:49:42   123 2002/02/12 19:57:45 4
     58     19       1828     39       2840    40 2002/02/12 19:52:39   117 2002/02/12 19:56:29 6
    123     74       5642     49       3484 Total
     2 Ip Version(s) found
3  Device Type Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Device Type
     42     42       2667      0          0     4 2002/02/12 19:49:48   123 2002/02/12 19:57:45 34(Loopback)
     23     13       1147     10        644     1 2002/02/12 19:49:42   103 2002/02/12 19:55:48 39(QDIO Ethernet)
      6      3        324      3        180    49 2002/02/12 19:52:58    54 2002/02/12 19:52:58 51(Loopback6)
     52     16       1504     36       2660    40 2002/02/12 19:52:39   117 2002/02/12 19:56:29 53(QDIO Ethernet6)
    123     74       5642     49       3484 Total
     4 Device Type(s) found
4  Interface Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Interface
     42     42       2667      0          0     4 2002/02/12 19:49:48   123 2002/02/12 19:57:45 LOOPBACK
      6      3        324      3        180    49 2002/02/12 19:52:58    54 2002/02/12 19:52:58 LOOPBACK
     23     13       1147     10        644     1 2002/02/12 19:49:42   103 2002/02/12 19:55:48 OSAQDIOL
     52     16       1504     36       2660    40 2002/02/12 19:52:39   117 2002/02/12 19:56:29 OSAQDIO4
    123     74       5642     49       3484 Total
     4 Interface(s) found
Interface Address Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Interface
      5      5        699      0          0     1 2002/02/12 19:49:42   103 2002/02/12 19:55:48 OSAQDIOLINK
                                                  Addr: 9.67.115.63
     42     42       2667      0          0     4 2002/02/12 19:49:48   123 2002/02/12 19:57:45 LOOPBACK
                                                  Addr: 9.67.115.5
     18      8        448     10        644    16 2002/02/12 19:51:17    33 2002/02/12 19:52:21 OSAQDIOLINK
                                                  Addr: 9.67.115.5
     31     14       1360     17       1340    40 2002/02/12 19:52:39   116 2002/02/12 19:56:29 OSAQDIO46
                                                  Addr: FEC9:C2D4::6:2900:EDC:217C
     21      2        144     19       1320    46 2002/02/12 19:52:44   117 2002/02/12 19:56:29 OSAQDIO46
                                                  Addr: FE80::6:2900:EDC:217C
      6      3        324      3        180    49 2002/02/12 19:52:58    54 2002/02/12 19:52:58 LOOPBACK6
                                                  Addr: FEC9:C2D4::9:67:115:5
    123     74       5642     49       3484 Total
     6 Interface Address(s) found
Asid Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Asid
     34      0          0     34       2440    16 2002/02/12 19:51:17   111 2002/02/12 19:56:24 002A
     89     74       5642     15       1044     1 2002/02/12 19:49:42   123 2002/02/12 19:57:45 002B
    123     74       5642     49       3484 Total
     2 Asid(s) found

5  Protocol Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Protocol
     30     29       1624      1        284     5 2002/02/12 19:49:48   123 2002/02/12 19:57:45 1(ICMP)
      4      2        120      2        160    99 2002/02/12 19:55:20   102 2002/02/12 19:55:46 6(TCP)
     53     26       2190     27       1440     1 2002/02/12 19:49:42   122 2002/02/12 19:57:45 17(UDP)
     36     17       1708     19       1600    41 2002/02/12 19:52:39   117 2002/02/12 19:56:29 58(ICMPv6)
    123     74       5642     49       3484 Total
     4 Protocol(s) found
6  IP Address Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss
      9      5        280      4        404    16 2002/02/12 19:51:17    27 2002/02/12 19:52:21
                                                  Addr: 9.67.115.1
     60     50       3115     10        644     4 2002/02/12 19:49:48   123 2002/02/12 19:57:45
                                                  Addr: 9.67.115.5
      5      5        699      0          0     1 2002/02/12 19:49:42   103 2002/02/12 19:55:48
                                                  Addr: 9.67.115.17
      5      5        699      0          0     1 2002/02/12 19:49:42   103 2002/02/12 19:55:48
                                                  Addr: 9.67.115.63
      9      3        168      6        240    22 2002/02/12 19:52:21    33 2002/02/12 19:52:21
                                                  Addr: 9.67.115.69
      4      0          0      4        240    55 2002/02/12 19:53:17    67 2002/02/12 19:53:32
                                                  Addr: FE80::6:2900:ADC:217C
     21      2        144     19       1320    46 2002/02/12 19:52:44   117 2002/02/12 19:56:29
                                                  Addr: FE80::6:2900:EDC:217C
      4      2        144      2        144    45 2002/02/12 19:52:44   105 2002/02/12 19:55:51
                                                  Addr: FE80::202:55FF:FE64:2DE7
      2      1         72      1         72   116 2002/02/12 19:56:29   117 2002/02/12 19:56:29
                                                  Addr: FE80::206:2AFF:FE66:C800
      6      3        216      3        216    76 2002/02/12 19:54:02    94 2002/02/12 19:55:09
                                                  Addr: FE80::206:2AFF:FE71:4400
     31     14       1360     17       1340    40 2002/02/12 19:52:39   116 2002/02/12 19:56:29
                                                  Addr: FEC9:C2D4::6:2900:EDC:217C
```

```
      6      3        324      3        180    49 2002/02/12 19:52:58    54 2002/02/12 19:52:58
                                                  Addr: FEC9:C2D4::9:67:115:5
      8      4        348      4        260    40 2002/02/12 19:52:39   102 2002/02/12 19:55:46
                                                  Addr: FEC9:C2D4::9:67:115:17
      3      2        376      1        304   110 2002/02/12 19:56:24   113 2002/02/12 19:56:24
                                                  Addr: FEC9:C2D4::206:2AFF:FE66:C800
      8      4        348      4        260    88 2002/02/12 19:55:04   100 2002/02/12 19:55:20
                                                  Addr: FEC9:C2D4::206:2AFF:FE71:4400
      5      0          0      5        300    74 2002/02/12 19:53:57    82 2002/02/12 19:54:17
                                                  Addr: FEC9:C2D4:1::9:67:114:44
      1      0          0      1         72    41 2002/02/12 19:52:39    41 2002/02/12 19:52:39
                                                  Addr: FF02::1:FF15:17
      1      0          0      1         72   111 2002/02/12 19:56:24   111 2002/02/12 19:56:24
                                                  Addr: FF02::1:FF66:C800
      1      0          0      1         72    89 2002/02/12 19:55:04    89 2002/02/12 19:55:04
                                                  Addr: FF02::1:FF71:4400
      9      0          0      9        648    56 2002/02/12 19:53:17    66 2002/02/12 19:53:31
                                                  Addr: FF02::1:FFDC:217C
    198    103       8293     95       6788 Total
    20 IP Address(s) found
Qos Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Qos
      5      5        280      0          0    18 2002/02/12 19:51:54    27 2002/02/12 19:52:21 6(Internetwork)
      1      1         60      0          0   100 2002/02/12 19:55:20   100 2002/02/12 19:55:20 96(Unknown)
      6      6        432      0          0    77 2002/02/12 19:54:02   116 2002/02/12 19:56:29 112(Unknown)
     12     12        772      0          0 Total
     3 Qos(s) found
 7  Tcp Port Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Tcp Port
      4      2        120      2        160    99 2002/02/12 19:55:20   102 2002/02/12 19:55:46 21(ftp)
      2      1         60      1         80    99 2002/02/12 19:55:20   100 2002/02/12 19:55:20 1026()
      2      1         60      1         80   101 2002/02/12 19:55:46   102 2002/02/12 19:55:46 1027()
      8      4        240      4        320 Total
     3 Tcp Port(s) found

 8  Udp Port Report
  Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Udp Port
      3      3        234      0          0     1 2002/02/12 19:49:42     3 2002/02/12 19:49:44 137(netbios-ns)
      2      2        465      0          0    83 2002/02/12 19:54:38   103 2002/02/12 19:55:48 138(netbios-dgm)
     21     21       1491      0          0     4 2002/02/12 19:49:48   122 2002/02/12 19:57:45 161(snmp)
      3      3        213      0          0     4 2002/02/12 19:49:48     8 2002/02/12 19:49:57 1035()
      3      3        213      0          0    10 2002/02/12 19:51:06    14 2002/02/12 19:51:15 1036()
      3      3        213      0          0    34 2002/02/12 19:52:24    38 2002/02/12 19:52:33 1037()
      3      3        213      0          0    68 2002/02/12 19:53:42    72 2002/02/12 19:53:51 1038()
      3      3        213      0          0    84 2002/02/12 19:55:00    97 2002/02/12 19:55:09 1039()
      3      3        213      0          0   106 2002/02/12 19:56:18   114 2002/02/12 19:56:27 1040()
      3      3        213      0          0   118 2002/02/12 19:57:36   122 2002/02/12 19:57:45 1041()
     27      0          0     27       1440    17 2002/02/12 19:51:54    95 2002/02/12 19:55:09 32810()
      7      0          0      7        380    17 2002/02/12 19:51:54    88 2002/02/12 19:55:04 33435()
      7      0          0      7        380    19 2002/02/12 19:51:54    92 2002/02/12 19:55:04 33436()
      7      0          0      7        380    20 2002/02/12 19:51:59    95 2002/02/12 19:55:09 33437()
      3      0          0      3        160    28 2002/02/12 19:52:21    81 2002/02/12 19:54:12 33438()
      2      0          0      2        100    30 2002/02/12 19:52:21    82 2002/02/12 19:54:17 33439()
      1      0          0      1         40    32 2002/02/12 19:52:21    32 2002/02/12 19:52:21 33440()
    101     47       3681     54       2880 Total
    17 Udp Port(s) found
Protocol Summary Report
             Input               Output              Total
Protocol Packets      Bytes  Packets      Bytes  Packets      Bytes
Tcp            2        120        2        160        4        280
Udp           26       2190       27       1440       53       3630
Icmp          46       3332       20       1884       66       5216
Other          0          0        0          0        0          0
Session Summary Report
 Input Output  First yyyy/mm/dd hh.mm.ss   Last yyyy/mm/dd hh.mm.ss Protocol
     5      1     16 2002/02/12 19:51:17     27 2002/02/12 19:52:21 ICMP     Lcl:      9.67.115.5-0
                                                                            Rmt:      9.67.115.1-0
     0      1     17 2002/02/12 19:51:54     17 2002/02/12 19:51:54 UDP      Lcl:      9.67.115.5-32810
                                                                            Rmt:      9.67.115.1-33435
     0      1     19 2002/02/12 19:51:54     19 2002/02/12 19:51:54 UDP      Lcl:      9.67.115.5-32810
                                                                            Rmt:      9.67.115.1-33436
     0      1     20 2002/02/12 19:51:59     20 2002/02/12 19:51:59 UDP      Lcl:      9.67.115.5-32810
                                                                            Rmt:      9.67.115.1-33437
    21      0      5 2002/02/12 19:49:48    123 2002/02/12 19:57:45 ICMP     Lcl:      9.67.115.5-0
                                                                            Rmt:      9.67.115.5-0
     3      0      4 2002/02/12 19:49:48      8 2002/02/12 19:49:57 UDP      Lcl:      9.67.115.5-161
                                                                            Rmt:      9.67.115.5-1035
     3      0     10 2002/02/12 19:51:06     14 2002/02/12 19:51:15 UDP      Lcl:      9.67.115.5-161
                                                                            Rmt:      9.67.115.5-1036
```

```
3      0     34 2002/02/12 19:52:24     38 2002/02/12 19:52:33 UDP      Lcl:      9.67.115.5-161
                                                                        Rmt:      9.67.115.5-1037
3      0     68 2002/02/12 19:53:42     72 2002/02/12 19:53:51 UDP      Lcl:      9.67.115.5-161
                                                                        Rmt:      9.67.115.5-1038
3      0     84 2002/02/12 19:55:00     97 2002/02/12 19:55:09 UDP      Lcl:      9.67.115.5-161
                                                                        Rmt:      9.67.115.5-1039
3      0    106 2002/02/12 19:56:18    114 2002/02/12 19:56:27 UDP      Lcl:      9.67.115.5-161
                                                                        Rmt:      9.67.115.5-1040
3      0    118 2002/02/12 19:57:36    122 2002/02/12 19:57:45 UDP      Lcl:      9.67.115.5-161
                                                                        Rmt:      9.67.115.5-1041
3      0     29 2002/02/12 19:52:21     33 2002/02/12 19:52:21 ICMP     Lcl:      9.67.115.5-0
                                                                        Rmt:     9.67.115.69-0
0      1     22 2002/02/12 19:52:21     22 2002/02/12 19:52:21 UDP      Lcl:      9.67.115.5-32810
                                                                        Rmt:     9.67.115.69-33435
0      1     24 2002/02/12 19:52:21     24 2002/02/12 19:52:21 UDP      Lcl:      9.67.115.5-32810
                                                                        Rmt:     9.67.115.69-33436
0      1     26 2002/02/12 19:52:21     26 2002/02/12 19:52:21 UDP      Lcl:      9.67.115.5-32810
                                                                        Rmt:     9.67.115.69-33437
0      1     28 2002/02/12 19:52:21     28 2002/02/12 19:52:21 UDP      Lcl:      9.67.115.5-32810
                                                                        Rmt:     9.67.115.69-33438
0      1     30 2002/02/12 19:52:21     30 2002/02/12 19:52:21 UDP      Lcl:      9.67.115.5-32810
                                                                        Rmt:     9.67.115.69-33439
0      1     32 2002/02/12 19:52:21     32 2002/02/12 19:52:21 UDP      Lcl:      9.67.115.5-32810
                                                                        Rmt:     9.67.115.69-33440
3      0      1 2002/02/12 19:49:42      3 2002/02/12 19:49:44 UDP      Lcl:     9.67.115.63-137
                                                                        Rmt:     9.67.115.17-137
2      0     83 2002/02/12 19:54:38    103 2002/02/12 19:55:48 UDP      Lcl:     9.67.115.63-138
                                                                        Rmt:     9.67.115.17-138

0      1     55 2002/02/12 19:53:17     55 2002/02/12 19:53:17 UDP      Lcl: FE80::6:2900:EDC:217C-32810
                                                                        Rmt: FE80::6:2900:ADC:217C-33435
0      1     59 2002/02/12 19:53:22     59 2002/02/12 19:53:22 UDP      Lcl: FE80::6:2900:EDC:217C-32810
                                                                        Rmt: FE80::6:2900:ADC:217C-33436
0      1     63 2002/02/12 19:53:27     63 2002/02/12 19:53:27 UDP      Lcl: FE80::6:2900:EDC:217C-32810
                                                                        Rmt: FE80::6:2900:ADC:217C-33437
0      1     67 2002/02/12 19:53:32     67 2002/02/12 19:53:32 UDP      Lcl: FE80::6:2900:EDC:217C-32810
                                                                        Rmt: FE80::6:2900:ADC:217C-33438
0      2     46 2002/02/12 19:52:44    105 2002/02/12 19:55:51 ICMPv6   Lcl: FE80::6:2900:EDC:217C-0
                                                                        Rmt: FE80::202:55FF:FE64:2DE7-0
0      1    117 2002/02/12 19:56:29    117 2002/02/12 19:56:29 ICMPv6   Lcl: FE80::6:2900:EDC:217C-0
                                                                        Rmt: FE80::206:2AFF:FE66:C800-0
2      3     76 2002/02/12 19:54:02     94 2002/02/12 19:55:09 ICMPv6   Lcl: FE80::6:2900:EDC:217C-0
                                                                        Rmt: FE80::206:2AFF:FE71:4400-0
0      9     56 2002/02/12 19:53:17     66 2002/02/12 19:53:31 ICMPv6   Lcl: FE80::6:2900:EDC:217C-0
                                                                        Rmt: FF02::1:FFDC:217C-0
2      0     45 2002/02/12 19:52:44    104 2002/02/12 19:55:51 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                        Rmt: FE80::202:55FF:FE64:2DE7-0
1      0    116 2002/02/12 19:56:29    116 2002/02/12 19:56:29 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                        Rmt: FE80::206:2AFF:FE66:C800-0
1      0     93 2002/02/12 19:55:09     93 2002/02/12 19:55:09 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                        Rmt: FE80::206:2AFF:FE71:4400-0
1      1    101 2002/02/12 19:55:46    102 2002/02/12 19:55:46 TCP      Lcl: FEC9:C2D4::6:2900:EDC:217C-1027
                                                                        Rmt: FEC9:C2D4::9:67:115:17-21
0      1     40 2002/02/12 19:52:39     40 2002/02/12 19:52:39 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4::9:67:115:17-33435
0      1     44 2002/02/12 19:52:39     44 2002/02/12 19:52:39 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4::9:67:115:17-33436
0      1     47 2002/02/12 19:52:44     47 2002/02/12 19:52:44 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4::9:67:115:17-33437
3      0     42 2002/02/12 19:52:39     48 2002/02/12 19:52:44 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                        Rmt: FEC9:C2D4::9:67:115:17-0
2      1    110 2002/02/12 19:56:24    113 2002/02/12 19:56:24 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                        Rmt: FEC9:C2D4::206:2AFF:FE66:C800-0
1      1     99 2002/02/12 19:55:20    100 2002/02/12 19:55:20 TCP      Lcl: FEC9:C2D4::6:2900:EDC:217C-1026
                                                                        Rmt: FEC9:C2D4::206:2AFF:FE71:4400-21
0      1     88 2002/02/12 19:55:04     88 2002/02/12 19:55:04 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4::206:2AFF:FE71:4400-33435
0      1     92 2002/02/12 19:55:04     92 2002/02/12 19:55:04 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4::206:2AFF:FE71:4400-33436
0      1     95 2002/02/12 19:55:09     95 2002/02/12 19:55:09 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4::206:2AFF:FE71:4400-33437
3      0     90 2002/02/12 19:55:04     96 2002/02/12 19:55:09 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                        Rmt: FEC9:C2D4::206:2AFF:FE71:4400-0
0      1     74 2002/02/12 19:53:57     74 2002/02/12 19:53:57 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4:1::9:67:114:44-33435
0      1     75 2002/02/12 19:54:02     75 2002/02/12 19:54:02 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4:1::9:67:114:44-33436
0      1     78 2002/02/12 19:54:07     78 2002/02/12 19:54:07 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                        Rmt: FEC9:C2D4:1::9:67:114:44-33437
```

```
     0      1     81 2002/02/12 19:54:12     81 2002/02/12 19:54:12 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                              Rmt: FEC9:C2D4:1::9:67:114:44-33438
     0      1     82 2002/02/12 19:54:17     82 2002/02/12 19:54:17 UDP      Lcl: FEC9:C2D4::6:2900:EDC:217C-32810
                                                                              Rmt: FEC9:C2D4:1::9:67:114:44-33439
     0      1     41 2002/02/12 19:52:39     41 2002/02/12 19:52:39 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                              Rmt: FF02::1:FF15:17-0
     0      1    111 2002/02/12 19:56:24    111 2002/02/12 19:56:24 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                              Rmt: FF02::1:FF66:C800-0
     0      1     89 2002/02/12 19:55:04     89 2002/02/12 19:55:04 ICMPv6   Lcl: FEC9:C2D4::6:2900:EDC:217C-0
                                                                              Rmt: FF02::1:FF71:4400-0
     0      1     49 2002/02/12 19:52:58     49 2002/02/12 19:52:58 UDP      Lcl: FEC9:C2D4::9:67:115:5-32810
                                                                              Rmt: FEC9:C2D4::9:67:115:5-33435
     0      1     51 2002/02/12 19:52:58     51 2002/02/12 19:52:58 UDP      Lcl: FEC9:C2D4::9:67:115:5-32810
                                                                              Rmt: FEC9:C2D4::9:67:115:5-33436
     0      1     53 2002/02/12 19:52:58     53 2002/02/12 19:52:58 UDP      Lcl: FEC9:C2D4::9:67:115:5-32810
                                                                              Rmt: FEC9:C2D4::9:67:115:5-33437
     3      0     50 2002/02/12 19:52:58     54 2002/02/12 19:52:58 ICMPv6   Lcl: FEC9:C2D4::9:67:115:5-0
                                                                              Rmt: FEC9:C2D4::9:67:115:5-0

 9     55 session(s) found
 10    123 records processed for this report
 11  Recording ended at   2002/02/12 19:57:45.836155
Recording started at 2002/02/12 19:49:42.788207
The duration was               00:08:03.047947
 12    304 is the maximum packet data length
 13  16384 bytes of storage used to create this report
   123 requests for 14992 bytes of storage were made
```

1

The standard statistics shown with all executions of the SYSTCPDA packet trace formatter.

ctrace records processed
The total number of CTRACE records given to the SYSTCPDA packet trace formatter by IPCS.

segmented trace records read
The total number of packets that spanned multiple CTRACE records.

segmented trace records were lost
The total number of segmented packets records that could not be put back together.

trace records read
The total number of complete trace records.

records could not be validated
The number of incomplete CTRACE records that could not be used.

records passed filtering
The number of records that were successfully formatted.

packet trace records processed
The number of records that were packet trace records.

data trace records processed
The number of records that were data trace records.

2

The totals by record type (Packet trace, X25, and data trace).

3

The totals by device type for packet trace records.

4

The totals by Interface or Link Name for packet trace records.

5

The totals by Protocol number for packet trace records.

6
The totals by IP Address. Both the destination and source IP addresses are counted except when they are the same with in a record.

7
The totals by TCP Port number. Both the destination and source port numbers are counted except when they are the same within a record.

8
The totals by UDP port number. Both the destination and source port numbers are counted except when they are the same within a record.

Restriction: Reports **2** through **8** are shown only when STATISTICS(DETAIL) is specified in the OPTIONS string.

9
The totals by session partner pairs (IP addresses, protocol number, and port numbers).

10
The number of records processed for the statistics report.

11
The time stamp of the first record in the input file, the time stamp of the last record in the input, and the duration from the first to last record.

Note: Records that have been abbreviated are not shown in this example. The number of records that were abbreviated and the maximum abbreviated size are shown. Also, the number and maximum size of the records that were not abbreviated are shown.

12
The size of the largest packet found in the input file.

13
The number of records processed for the statistics report, the number of 1-KB blocks of storage required for this report, the number of storage requests, and the total amount of storage required for the requests.

The report by Jobname for data trace records is not shown. Each category of totals is broken down by:

- The total number of records
- The total number of inbound records
- The total amount of inbound data
- The total number of outbound records
- The total amount of outbound data
- The record number of the first record
- The time stamp of the first record
- The record number of the last record
- The time stamp of the last record

STREAM:
Purpose

Sometimes messages span multiple packets. TELNET and DNS are examples. The STREAM report (with the DUMP or FORMAT keywords) captures the entire stream of data.

Format

```
CTRACE COMP(SYSTCPDA) SUB((TCPCS)) SHORT OPTIONS((STREAM DUMP ASCII))
```

```
================================================================================
1  Streams Report
   60 Streams found
 23600 bytes of storage for the session report was allocated
1146880 bytes of storage for buffers was allocated

--------------------------------------------------------------------------------
2  Session: FEC9:C2D4::206:2AFF:FE71:4400-0 FEC9:C2D4::6:2900:EDC:217C-0 ICMPv6
 From: 2002/02/12 19:55:04.615118 to: 2002/02/12 19:55:09.636234
      3 packets found
 Stream buffer at 0A818000 for 20480 bytes. 144 bytes were used
      3 packets moved for 144 bytes
I - Inbound packet
O - Outbound packet

3  D Rcd #            Time           Delta      Seq #   Position Length    End_Pos
I    90 19:55:04.615118 00:00:00.000000         0          0     24         24
000000 FEC9C2D4 00000000 02062AFF FE714400  02010006 2A714400                  |..........*..qD.....*qD.        |
I    91 19:55:04.631206 00:00:00.016087        24         24     60         84
000000                                                 60000000 00141101 |                          `.......|
000020 FEC9C2D4 00000000 00062900 0EDC217C  FEC9C2D4 00000000 02062AFF FE714400 |..........)...!|..........*..qD.|
000040 802A829B 0014A3B6 01010000 3C697318  00095CAB                           |.*..........<is.....
```

1

After all the records have processed, the number of streams and the amount of storage required for the report and buffers are shown.

2

Each session is identified by the IP addresses, port number, and protocol. The time stamps of the first and last packet are shown along with the number of packets, the address, and size of the stream buffer.

3

When a stream is dumped, each packet and the data from the packet is shown. If there are gaps in the stream, the number of bytes skipped is displayed. The data about each packet formatted are:

D The direction of the packet: I for inbound and O for outbound.

Rcd #
The record number.

Time
The time stamp of the record.

Delta
The time from the first record of the stream.

Seq #
The sequence number of the TCP packet. For other packets it is the relative offset of the packet from the first packet.

Position
The relative offset of the packet.

Length
The number of bytes in the packet.

End_Pos
The ending sequence number.

Formatting packet trace using a batch job

A Packet Trace can also be formatted through the use of a batch job. The following example is an example of JCL for a batch job:

```
//jobname   DD (accounting),pgmname,CLASS=A,MSGCLASS=A
//DUMP    EXEC PGM=IKJEFT01
//STEPLIB   DD DISP=SHR,DSN=hlq.MIGLIB
//SYSPRINT  DD SYSOUT=*
```

```
//SYSUDUMP  DD SYSOUT=*
//SYSTSPRT  DD SYSOUT=*
//PRINTER   DD SYSOUT=*
//SYSPROC   DD DISP=SHR,DSN=SYS1.CLIST
//          DD DISP=SHR,DSN=SYS1.SBLSCLI0
//IPCSPARM DD DISP=SHR,DSN=SYS1.PARMLIB
//         DD DISP=SHR,DSN=CPAC.PARMLIB
//         DD DISP=SHR,DSN=SYS1.IBM.PARMLIB
//IPCSPRNT  DD SYSOUT=*
//IPCSTOC   DD SYSOUT=*
//IPCSDDIR  DD DISP=SHR,DSN=userid.IPCS.DMPDIR
//SYSTSIN   DD *
 IPCS NOPARM
 SETDEF DA('ctrace.dataset')
 CTRACE COMP(SYSTCPDA) OPTIONS((systcpda_options_string))
 END
/*
```

Data trace (SYSTCPDA) for TCP/IP stacks

Use the DATTRACE command to trace socket data (transforms) into and out of the physical file structure (PFS). For TCP and UDP sockets, data trace also creates the following records:

- A Start record with State field API Data Flow Starts that indicates the first data sent or received by the application for the associated socket
- An End record with State field API Data Flow Ends that indicates that the socket has been closed

DATTRACE operates with the following APIs:

- REXX
- C-sockets
- IMS
- CICS
- Native z/OS UNIX
- Macro
- CALL Instruction

See z/OS Communications Server: IP System Administrator's Commands for information about the format of the data trace command (VARY DATTRACE).

Starting data trace

You can start data trace for all job names using the VARY command:

```
V TCPIP,tcpprocname,DAT
```

Tips:

- To use any VARY command, the user must be authorized in RACF. This replaces the old OBEY list authorization.
- Each user's RACF profile must have access for a resource of the form MVS.VARY.TCPIP.*xxx*, where *xxx* is the first eight characters of the command name. For data trace, this would be MVS.VARY.TCPIP.DATTRACE.
- Traces are placed in an internal buffer, which can then be written out using an external writer. The MVS TRACE command must also be issued for component SYSTCPDA to activate the data trace.

Displaying data traces

You can use the Netstat CONFIG/-f command to display data traces. Figure 21 shows a data trace for a single entry.

```
netstat -p TCPCS -f
...
Data Trace Setting:
JobName: *           TrRecCnt: 00000006  Length: FULL
IpAddr:  *                 SubNet: *
PortNum: *
```

Figure 21. Data trace: Single entry

Figure 22 shows a data trace for multiple entries.

```
netstat -p TCPCS -f
...
Data Trace Setting:
JobName: MEGA4       TrRecCnt: 00000000  Length: FULL
IpAddr:  127.0.0.3         SubNet: *
PortNum: *
JobName: *           TrRecCnt: 00000000  Length: FULL
IpAddr:  127.0.0.9         SubNet: *
PortNum: *
```

Figure 22. Data trace with multiple entries

Formatting data traces using IPCS

Data trace records are written to the same CTRACE component as packet trace records (SYSTCPDA). Thus, all the IPCS formatting features for packet trace are also available for data trace. You can use the ENTIDLIST parameter to isolate data trace records and packet trace records from each other. For an example of data trace records formatted by the FORMAT option, see the sample below. For a full description of the FORMAT option and information about other options that can be used to format data trace records, see "Formatting packet traces using IPCS" on page 96.

SYSTCPDA can create two types of records:

Common data trace records
SYSTCPDA creates common data trace records for every data exchange between the two endpoints of a session.

Start and End records
For TCP and UDP sockets, SYSTCPDA creates start and end records to delineate the logical boundaries of a session[1]. The start and end records are created for the following socket processing:

- Initial socket read or write
- Close of a socket

The main difference between the two types of records is the absence of actual data in start and end records. Common data trace records have a field called Data, which contain payload data that is transmitted between the two endpoints. Start and end records also contain a field called State, which indicates whether the record represents the start or end of the session.

1. The term session here should be read as a logical connection between two endpoints, independent of the protocol used (which can be TCP or UDP).

Figure 23 is an example of the FORMAT option output, which shows a set of records for a socket, including the start data trace record, a common data trace record, and the end data trace record.

```
 63 MVS182   DATA     00000005 13:18:52.705250 Data Trace
1 To Jobname        : USER12                                  Full=0
2  Tod Clock        : 2009/06/22 13:18:52.705248              Cid: 00000058
3  Domain           : AF_Inet          Type: Stream           Protocol: TCP
4  State            : API Data Flow Starts
   Segment #        : 0                Flags:  Out
5  Source           : 10.81.2.5
6  Destination      : 10.81.2.1
   Source Port      : 0           Dest Port: 2000  Asid: 0041 TCB: 006FF1D8

-------------------------------------------------------------------------------
   64 MVS182   DATA     00000005 13:18:52.705252 Data Trace
 To Jobname     : USER12                                  Full=39
  Tod Clock     : 2009/06/22 13:18:52.705252              Cid: 00000058
  Domain        : AF_Inet          Type: Stream           Protocol: TCP
  Segment #     : 0                Flags: Adj Out
  Source        : 10.81.2.5
  Destination   : 10.81.2.1
  Source Port   : 2000             Dest Port: 0     Asid: 0041 TCB: 006FF1D8

7  Data            : 39     Data Length: 39                  Offset: 0
  000000 E3D5E3D6 F0F0F0F5 40C9D7A5 F440E4C4 D740F1A2 A340E6D9 C9E3C540 C6D9D6D4
  000020 40

-------------------------------------------------------------------------------
   63 MVS182   DATA     00000005 13:18:52.705250 Data Trace
 To Jobname     : USER12                                  Full=0
  Tod Clock     : 2009/06/22 13:18:52.705254              Cid: 00000058
  Domain        : AF_Inet          Type: Stream           Protocol: TCP
  State         : API Data Flow Ends
  Segment #     : 0                Flags: None
  Source        : 10.81.2.5
  Destination   : 10.81.2.1
  Source Port   : 0           Dest Port: 2000 Asid: 0041 TCB: 006FF1D8
```

Figure 23. FORMAT option output

1 *jobname*
: Provides the name of the job that performed the socket read or write operation.

Full or Abbrev
: Indicates whether the whole packet was traced and provides the length of the traced data.

- The value Full = x indicates that the whole packet was traced.
- The value Abbrev = x indicates that a truncated portion of the packet was traced.

2 Tod Clock
: Provides the timestamp of the time when the read or write operation took place.

Cid: *xxxxxxxx*
: Provides the connection id that uniquely identifies the session between the two endpoints. As shown in the example, the Cid is the same for the three records.

3 Domain
: Indicates whether this is IPv4 (AF_Inet) or IPv6 (AF_Inet6).

Type and Protocol
Provide the type of traffic (e.g. Stream, Datagram, Raw) and protocol used.

4 State
This field is only displayed for a start or end record and indicates whether this is a start or an end record.

5 Source and Destination
Provide the source and destination IP addresses of the two end points. For TCP sockets, these values are always provided. For UDP sockets, these values are provided only when a connect or bind function was executed against the socket.

6 Source Port and Dest Port
Provide the port numbers of the source and destination addresses.

Asid and TCB
Provide the address space id and TCB address.

7 Data
Provides the length of the payload data and is followed by a hexadecimal display of the actual payload data.

Intrusion Detection Services trace (SYSTCPIS)

When starting the TCP stack, the stack reads the CTIIDS00 parmlib member to determine the size to reserve for the SYSTCPIS CTRACE. You can override this default by starting TCP/IP with the PARM option and the keyword IDS=*xx*, where *xx* is the suffix of the CTIIDS*xx* PARMLIB member. In the following example, the trace searches for PARMLIB member CTIIDSA3.

```
S tcpiproc,PARM='IDS=A3'
```

If the parmlib member is not found or the member contains data that is not valid, the following message is displayed.

```
EZZ4210I CTRACE DEFINE FAILED FOR CTIIDS00
```

If the EZZ4210I message indicates the parmlib member name CTIIDS00, the IDS CTRACE space is set up using the default BUFSIZE of 32M.

The CTIIDS00 member is used to specify the IDS CTRACE parameters. To eliminate this message, ensure that a CTIIDS00 member exists within Parmlib and that the options are correctly specified. A sample CTIIDS00 member is shipped with z/OS Communications Server.

Packets are traced based on IDS policy. See Intrusion Detection Services in z/OS Communications Server: IP Configuration Guide for information about defining policy.

See Chapter 27, “Diagnosing intrusion detection problems,” on page 699 for additional information about diagnosing policy problems.

```
/*******************************************************************/
/*                                                                 */
/*  IBM Communications Server for z/OS                             */
/*  SMP/E Distribution Name: CTIIDS00                              */
/*                                                                 */
/*  MEMBER:  CTIIDS00                                              */
/*                                                                 */
/*                                                                 */
/*  Copyright:                                                     */
/*               Licensed Materials - Property of IBM              */
/*               5694-A01                                          */
/*               (C) Copyright IBM Corp. 2001, 2003                */
/*                                                                 */
/*                                                                 */
/*  Status:      CSV1R5                                            */
/*                                                                 */
/*                                                                 */
/*  DESCRIPTION = This parmlib member causes IDS component trace   */
/*               for the TCP/IP product to be initialized with a   */
/*               trace buffer size of 32M.                         */
/*                                                                 */
/*               This parmlib members only lists those TRACEOPTS   */
/*               value specific to TCP/IP.  For a complete list    */
/*               of TRACEOPTS keywords and their values see        */
/*               z/OS MVS INITIALIZATION AND TUNING REFERENCE.     */
/*                                                                 */
/* $MAC(CTIIDS00) PROD(TCPIP): Component Trace SYS1.PARMLIB member */
/*                                                                 */
/*                                                                 */
/*******************************************************************/
 TRACEOPTS
 /* ---------------------------------------------------------------- */
 /*   ON OR OFF: PICK 1                                              */
 /* ---------------------------------------------------------------- */
           ON
 /*        OFF                                                      */
 /* ---------------------------------------------------------------- */
 /*   BUFSIZE: A VALUE IN RANGE 16M TO 256M                          */
 /* ---------------------------------------------------------------- */
           BUFSIZE(32M)
 /*        WTR(wtr_procedure) WRAP|NOWRAP                           */
```

Figure 24. SYS1.PARMLIB member CTIIDS00

Restrictions

For IDS trace records the COMP keyword must be SYSTCPIS. Because there are no EXCEPTION records for IDS trace, the EXCEPTION keyword must not be specified.

CTRACE keywords on SYSTCPIS

The following describes those CTRACE keywords that affect SYSTCPIS processing.

ENTIDLIST
Use the ENTIDLIST keyword to select trace records with a specific ProbeId.

JOBLIST, JOBNAME
Use the JOBLIST and JOBNAME keywords to select trace records with a matching job name. Also, use the JOBNAME keyword in the OPTIONS list to select records.

ASIDLIST
Use the ASIDLIST to select trace records with a matching Asid.

GMT
The time stamps are converted to GMT time.

LOCAL
The time stamps are converted to LOCAL time.

SHORT
If the OPTIONS keyword does not specify any reports, format the trace records. Equivalent to the FORMAT option.

FULL
If the OPTIONS keyword does not specify any reports, format and dump the trace records. Equivalent to the FORMAT and DUMP options.

SUMMARY
If the OPTIONS keyword does not specify any reports, create a one line summary for each trace record. Equivalent to the SUMMARY option.

TALLY
If the OPTIONS keyword does not specify any reports, then count the trace records. Equivalent to the STATISTICS option.

START and STOP
These keywords limit the trace records seen by the packet trace formatter. The START keyword determines the time when records are seen by the packet trace report formatter. The STOP keyword determines the time when records are no longer seen by the packet trace report formatter.

Rule: CTRACE always uses the time the trace record was moved to the buffer for START and STOP times.

LIMIT
Determines the number of records that the packet trace formatter is allowed to process. See the RECORDS keyword value in OPTIONS.

USEREXIT
The CTRACE USEREXIT is not called because the packet trace formatter tells CTRACE to skip all the records. Therefore, the packet trace formatter calls the CTRACE USEREXIT before testing the records with the filtering criteria. If it returns a nonzero return code, the record is skipped. The USEREXIT can also be used in the OPTIONS keyword. It is called after the record has met all the filtering criteria in the OPTIONS keyword. The OPTIONS keyword provides a means of entering additional keywords for record selection and formatting.

SYSTCPIS OPTIONS

Use the CTRACE OPTIONS string to enter more keywords for record selection and formatting intrusion detection services (IDS) traces (COMP=SYSTCPIS). See "Syntax" on page 55 for the complete syntax of CTRACE.

The syntax for the OPTIONS keyword for IDS trace is the same as the syntax for the packet trace OPTIONS keyword, with a few exceptions. See "OPTIONS syntax" on page 97 for the values for the OPTIONS keyword.

You can use the following additional keywords to select records in an IDS trace:

Record Identifiers:

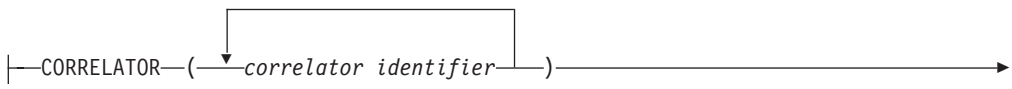

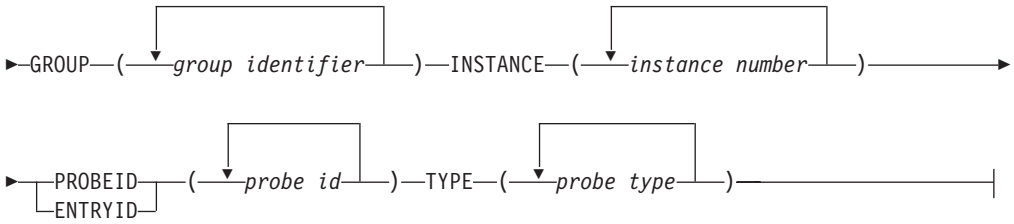

The following keywords that are available for packet trace are not available for selecting records in an IDS trace: ARP, DATTRACE, DEVICEID, ETHTYPE, HPRDIAG, MACADDR, PACKETTRACE, PKTTRACE, VLANID, and X25.

OPTIONS keywords

The following are keywords used for the OPTIONS component routine parameters.

CORRELATOR
Select trace records with one of the matching correlator identifiers. Up to 16 identifiers can be specified. Each identifier in the list can be a range: low_number:high_number. Values can be decimal (nnnnnnnnnn) or hexadecimal (X'hhhhhhhh'). This filter associates packets in the trace with an IDS event message in syslogd or the system console.

GROUP
Select trace records with one of the matching group identifiers. The following group identifiers can be specified:

- TCPTR
- UDPTR
- SCAN
- ATTACK

INSTANCE
Select trace records with one of the matching instance identifiers. The identifiers can be in decimal (nnnnn) or hexadecimal (x'hhhhhhhh'). The instance identifier is the lower 2 bytes of the probe identifier. Up to 16 identifiers can be specified.

PROBEID
Select trace records with one of the matching probe identifiers. The identifiers can be expressed in decimal (nnnnn) or hexadecimal (x'hhhhhhhh'). Up to 16 identifiers can be specified. You can also specify the probe identifiers on the ENTIDLIST keyword of the CTRACE subcommand. See z/OS Communications Server: IP and SNA Codes for additional information about probe identifiers.

TYPE(probe type identifier)
Select trace records with one of the matching probe type identifiers. The *probe type identifier* is the first 2 bytes of the probe identifier. Up to 16 identifiers can be specified. You can use the following probe types:

Identifier
Type

0100
TCPTR

0200
UDPTR

0301
VSSCAN

0302
PSSCAN

0303
NORMSCAN

0401
MALFORMED

0402
RAW

0403
IPFRAGMENT

0404
ICMP

0405
IPOPT

0406
IPPROTO

0407
FLOOD

0408
PERPECHO

0409
DATAHIDE

040A
TCPQUESZ

040B
TCPSTALL

040C
RAW6

040D
NEXTHDRS

040E
DESTOPTS

040F
HOPOPTS

0410
EELDLCCHECK

0411
EEMALFORMED

0412
EEPORTCHECK

0413
EEXIDFLOOD

IDS reports

The SYSTCPIS CTRACE formatter is based on the SYSTCPDA formatter (and in fact shares many of the data structures and format routines) and includes the reports for the SYSTCPDA formatter. However, the REASSEMBLY, SESSION, and STREAMS reports might prove of little value for the SYSTCPIS, because they depend on having a more complete set of packets.

- The STATISTICS report (both SUMMARY and DETAIL) provides an overview of the data collected.
- The SUMMARY report provides one line per IDS event.
- The FORMAT, and DUMP reports format individual packets.
- The EXPORT and SNIFFER options write the packet to an external file for later analysis.

The following topics describe the various reports available.

OPTION
Purpose

List the selected options and default keyword values.

Format

The following command was used to obtain the example of this report.

```
CTRACE COMP(SYSTCPIS) SUB((TCPCS)) DSN('IBMUSER.CTRACE1')
OPTION((OPT SESS FORM))
REPORT
```

```
COMPONENT TRACE SHORT FORMAT
  SYSNAME(MVS118)
  COMP(SYSTCPIS)SUBNAME((TCPCS))
  OPTIONS((OPT SESS FORM))
1   DSNAME('IBMUSER.CTRACE1')

2  OPTIONS((Both Bootp(67,68) Cleanup(500) DelayAck(200,200) Domain(53)
  Finger(79) Flags() Format(Detail) Ftp(20,21) Gain(125,250) Gopher(70)
  Limit(999999999) Ntp(123) Option Noreassembly Router(520) Rpc(111) Segment
  Session(Detail) Smtp(25) Snmp(161,162) Speed(10,10) Telnet(23) Tftp(69)
  Time(37) Userexit() Www(80)
   ))
```

The following describes numbered areas of the example.

1
DSNAME is the name of the source data.

2
OPTIONS((...)) is a listing of the active options with default values.

SUMMARY
Purpose

Show one line of information about each record in the trace.

Format

The following command was used to obtain the example of this report.

```
CTRACE COMP(SYSTCPIS) SUB((TCPCS)) SUMMARY DSN('IBMUSER.CTRACE1')
```

```
 COMPONENT TRACE SUMMARY FORMAT

 SYSNAME(MVS118)

 COMP(SYSTCPIS)SUBNAME((TCPCS))

 DSNAME('IBMUSER.CTRACE1')

**** 2002/11/20
I - Inbound packet

O - Outbound packet

DP   Nr hh:mm:ss.mmmmmm IpId Group    Probe Id Corelatr JobName  Cid      DatLn Data Source/Destination
II 4521 17:38:32.175560 0000 SCAN     03030000       10 TCPCS    00000000    12 ICMP  9.42.105.71
                                                                                     9.42.104.38
IT 4522 17:38:45.130339 163F SCAN     03030026       11 FTPD1    00000020     0 TCP   9.2.197.34-46911
                                                                                     9.42.104.38-21
IT 4523 17:38:45.153474 173F SCAN     03030026       12 FTPD1    00000020     0 TCP   9.224.157.220-47167
                                                                                     9.42.104.38-21
IT 4524 17:38:45.170441 183F SCAN     03030026       13 FTPD1    00000020     0 TCP   9.74.208.131-47423
                                                                                     9.42.104.38-21
IT 4525 17:38:45.190606 193F SCAN     03030026       14 FTPD1    00000020     0 TCP   9.79.235.253-47679
                                                                                     9.42.104.38-21
IT 4526 17:38:45.213117 1A3F SCAN     03030026       15 FTPD1    00000020     0 TCP   9.40.107.43-47935
                                                                                     9.42.104.38-21
IT 5671 17:59:32.787165 0B3B ATTACK   04070002      277 FTPD1    00000020     0 TCP    9.42.104.38-21
                                                                                     9.84.160.95-47938
IT 5672 17:59:32.806700 0B1A ATTACK   04070002      277 FTPD1    00000020     0 TCP    9.42.104.38-21
                                                                                     9.156.214.250-44610
IT 5673 17:59:32.827193 0B1B ATTACK   04070002      277 FTPD1    00000020     0 TCP    9.42.104.38-21
                                                                                     9.150.148.96-44866
IT 5674 17:59:32.847730 0B1C ATTACK   04070002      277 FTPD1    00000020     0 TCP    9.42.104.38-21
                                                                                     9.48.42.177-45122
.
.
.

================================================================================
SYSTCPIS Trace Statistics
 2,583 ctrace records processed
     0 segmented trace records read
     0 segmented trace records were lost
 2,583 trace records read
     0 records could not be validated
 2,583 records passed filtering
 2,583 packet trace records processed
     0 data trace records processed
```

The following describe areas of the example.

D Direction of the packet:

- **I** Inbound packet
- **O** Outbound packet

P The packet protocol:

- **T** TCP
- **U** UDP
- **I** ICMP
- **G** IGMP
- **P** Other

Nr The CTRACE record number.

hh:mm:ss.mmmmmmm
The time stamp of the record.

IpId
The packet ID number in hexadecimal.

Group
The group assigned to the trace record. The value can be ATTACK, SCAN, UDPTR or TCPTR.

Probe Id
The probe identifier assigned to the trace record.

Corelatr
The correlator assigned to the trace record. Use this to correlate the trace data with console or syslog messages.

JobName
The job name assigned to the trace record.

Cid
The connection identifier assigned to the trace record.

DatLn
The length of the data.

Data
The protocol in the IP header.

Source/Destination
The source and destination IP address and port number.

FORMAT
Purpose

Format the CTRACE record header, the IP packet header, the protocol header, and the packet data. If one of the ports is a well-known port number and the SYSTCPIS supports data for the port number, the packet data is shown.

Format

The following command was used to obtain the example of this report.

```
CTRACE COMP(SYSTCPIS) SUB((TCPCS)) SHORT DSN('IBMUSER.CTRACE1')
OPTIONS((OPT FORMAT))
```

```
 COMPONENT TRACE SHORT FORMAT
 SYSNAME(MVS118)
 COMP(SYSTCPIS)SUBNAME((TCPCS))
 OPTIONS((OPT FORMAT))
 DSNAME('IBMUSER.CTRACE1')

OPTIONS((Both Bootp(67,68) Cleanup(500) DelayAck(200,200) Domain(53)
 Finger(79) Flags() Format(Detail) Ftp(20,21) Gain(125,250) Gopher(70)
 Limit(999999999) Gmt Ntp(123) Option Noreassembly Router(520) Rpc(111)
 Segment Smtp(25) Snmp(161,162) Speed(10,10) Telnet(23) Tftp(69) Time(37)
 Userexit() Www(80)
  ))

1  **** 2002/11/20
 RcdNr Sysname  Mnemonic Entry Id   Time Stamp    Description
 ----- -------- -------- -------- --------------- --------------------------------

 ------------------------------------------------------------------------------
```

```
2   4521 MVS118   SCAN     03030000 17:38:32.175560 Scan-Normal packet
3   From Link         : ETH1             Device: LCS Ethernet    Full=40
   Tod Clock        : 2002/11/20 17:38:32.175559              Module: EZBIPICM
   Job Name         : TCPCS            Asid: 01F7             Tcb: 00000000
   Cid              : 00000000         Correlator: 10
   Policy           : ScanEventIcmp-rule
4   IpHeader: Version : 4               Header Length: 20
   Tos              : 00               QOS: Routine Normal Service
   Packet Length    : 40               ID Number: 0000
   Fragment         : DontFragment     Offset: 0
   TTL              : 62               Protocol: ICMP           CheckSum: 5914 FFFF
   Source           : 9.42.105.71
   Destination      : 9.42.104.38

5   ICMP
   Type/Code        : ECHO             CheckSum: 5592 FFFF
   Id               : 0B3F             Seq: 0
6    Echo Data        : 12
000000 AEBCDB3D 03340A00 00000000                                  |...=.4......|

-------------------------------------------------------------------------------
 4522 MVS118   SCAN     03030026 17:38:45.130339 Scan Normal-TCP SYN  dropped
 From Link         : UNKNOWN          Device: Unknown:0       Full=40
  Tod Clock        : 2002/11/20 17:38:45.130338              Module: EZBTCPCN
  Job Name         : FTPD1            Asid: 01F7             Tcb: 00000000
  Cid              : 00000020         Correlator: 11
  Policy           : ScanEventHigh-rule
 IpHeader: Version : 4                Header Length: 20
  Tos              : 00               QOS: Routine Normal Service
  Packet Length    : 40               ID Number: 163F
  Fragment         :                  Offset: 0
  TTL              : 253              Protocol: TCP            CheckSum: 681C FFFF
  Source           : 9.2.197.34
  Destination      : 9.42.104.38

 TCP
  Source Port      : 46911 ()       Destination Port: 21    (ftp)
  Sequence Number  : 2397868413     Ack Number: 0
  Header Length    : 20             Flags: Syn
  Window Size      : 242            CheckSum: 4E53 B695 Urgent Data Pointer: 0000
.
.
.
 ===============================================================================
 SYSTCPIS Trace Statistics
  2,623 ctrace records processed
      0 segmented trace records read
      0 segmented trace records were lost
  2,623 trace records read
      0 records could not be validated
  2,623 records passed filtering
  2,623 packet trace records processed
      0 data trace records processed
```

The following describes numbered areas of the example.

1

The date of the trace records.

2

A summary line indicating the source of the trace record showing:

- The record number.
- The system name.
- The group name.
- The probe ID value (in hexadecimal).

- The time the record was moved to the trace buffer, or with the TOD option the time the trace data was captured.
- The description of the IDS event associated with the probe.

3

The trace header with these fields:

- The direction of the trace record: *From* or *To*.
- The link name.
- The device type.
- *Full* or *Abbrev* with amount of trace data available.
- The time the trace record was captured.
- The module that triggered the probe.
- The job name associated when the probe was triggered.
- The ASID of the address space when the probe was triggered.
- The system tcb pointer when the probe was triggered (or zero if in SRB mode).
- The CID (communications ID) of the session.
- The Event identifier, the upper 2 bytes of the PROBEID.
- The Correlator identifier.
- The name of the current policy. This might be the policy that triggered the probe or the name of the policy the session was using at the time the probe was triggered.

4

The IP header showing fields from the IPv4 4 header. The header length is the number of bytes for the header. The offset field is the number of bytes from the end of the IP header where the fragment appears. With the REASSEMBLY option active, this field always displays zeros.

5

The protocol header. In this example, it is an ICMP header.

6

Depending on the port number, the trace data might be formatted.

Guideline: If possible, the check sum of the packet is calculated. If the calculated value is X'FFFF', then the check sum is correct. If the calculated value is X'0000', then the check sum could not be calculated. The packet was incomplete or fragmented. Other values indicate a check sum error.

Using the protocol numbers and the well known port numbers, format routines are invoked to format standard packet data records. The port number for the PORT keywords define the port numbers to be used to invoke a format routine.

Port
Keyword

67, 68
BOOTP

67, 68
DHCP

53 Domain

79 Finger

70 Gopher

520
Rip

520
Router

111
RFC

25 SMTP

23 TELNET

69 TFTP

37 TIME

DUMP
Purpose

Format the IP header, protocol header and packet data in hexadecimal. The data can also be translated into EBCDIC, ASCII or both.

Format

The following command was used to obtain the example of this report.

```
CTRACE COMP(SYSTCPIS) SUB((TCPCS)) DSNAME('IBMUSER.CTRACE1') SHORT OPTIONS((OPT DUMP))
COMPONENT TRACE SHORT FORMAT
  SYSNAME(MVS118)
  COMP(SYSTCPIS)SUBNAME((TCPCS))
  OPTIONS((OPT DUMP))
  DSNAME('IBMUSER.CTRACE1')

 OPTIONS((Both Bootp(67,68) Cleanup(500) DelayAck(200,200) Domain(53)
  Dump(65535) Finger(79) Flags() Ftp(20,21) Gain(125,250) Gopher(70)
  Limit(999999999) Gmt Ntp(123) Option Noreassembly Router(520) Rpc(111)
  Segment Smtp(25) Snmp(161,162) Speed(10,10) Telnet(23) Tftp(69) Time(37)
  Userexit() Www(80)
   ))

 **** 2002/11/20
 RcdNr Sysname  Mnemonic Entry Id   Time Stamp    Description
 ----- -------- -------- -------- --------------- ------------------------------

 -------------------------------------------------------------------------------
  4521 MVS118   SCAN     03030000 17:38:32.175560 Scan-Normal packet
  From Link         : ETH1             Device: LCS Ethernet    Full=40
   Tod Clock        : 2002/11/20 17:38:32.175559               Module: EZBIPICM
   Job Name         : TCPCS            Asid: 01F7              Tcb: 00000000
   Cid              : 00000000         Correlator: 10
   Policy           : ScanEventIcmp-rule

1 IP Header          : 20
 000000 45000028 00004000 3E015914 092A6947  092A6826

2 Protocol Header    : 8
 000000 08005592 0B3F0000

3 Data               : 12     Data Length: 12
 000000 AEBCDB3D 03340A00 00000000          |.............     ...=.4......    |

 -------------------------------------------------------------------------------
  4522 MVS118   SCAN     03030026 17:38:45.130339 Scan Normal-TCP SYN  dropped
```

```
 From Link         : UNKNOWN          Device: Unknown:0       Full=40
  Tod Clock        : 2002/11/20 17:38:45.130338               Module: EZBTCPCN
  Job Name         : FTPD1            Asid: 01F7              Tcb: 00000000
  Cid              : 00000020         Correlator: 11
  Policy           : ScanEventHigh-rule

IP Header          : 20
000000 45000028 163F0000 FD06681C 0902C522  092A6826

Protocol Header    : 20
000000 B73F0015 8EEC917D 00000000 500200F2  4E530000

-------------------------------------------------------------------------------
 4523 MVS118   SCAN     03030026 17:38:45.153474 Scan Normal-TCP SYN  dropped
 From Link         : UNKNOWN          Device: Unknown:0       Full=40
  Tod Clock        : 2002/11/20 17:38:45.153473               Module: EZBTCPCN
  Job Name         : FTPD1            Asid: 01F7              Tcb: 00000000
  Cid              : 00000020         Correlator: 12
  Policy           : ScanEventHigh-rule
IP Header          : 20
000000 45000028 173F0000 FD068D84 09E09DDC  092A6826

Protocol Header    : 20
000000 B83F0015 76399A57 00000000 500200F2  5D2C0000
.
.
.

===============================================================================
SYSTCPIS Trace Statistics
 2,623 ctrace records processed
     0 segmented trace records read
     0 segmented trace records were lost
 2,623 trace records read
     0 records could not be validated
 2,623 records passed filtering
 2,623 packet trace records processed
     0 data trace records processed
```

The following describes numbered areas of the example.

1 The IP header is dumped with no translation.

2 The protocol header is dumped with no translation.

3 The packet data is dumped with the translation specified by the ASCII, BOTH, EBCDIC, or HEX keyword. The default is BOTH. The amount of data dumped can be limited by the value specified with the DUMP keyword. The default is 65535 bytes.

SNIFFER
Purpose

This report shows information that is written to the SNIFFER data set.

Restriction: Do not use the SNIFFER option when the CTRACE subcommands are used with the IPCS MERGE subcommand. The SNIFFER data file is written over by the multiple CTRACE commands that specify the SNIFFER option.

Format

The following command was used to obtain the example of this report.

```
ALLOC F(SNIFFER) DATASET(SNIFFER.TRC) LRECL(1600) RECFM(V B) +
REUSE TRACK SPACE(15 15)

CTRACE COMP(SYSTCPIS)  DSN('MWS.PQ33208.PTRACE4')+
OPTION((OPT  TALLY SNIFFER NOREASSEMBLY))
COMPONENT TRACE SHORT FORMAT
  SYSNAME(MVS142)
  COMP(SYSTCPIS)
  OPTIONS(( OPT  TALLY SNIFFER(4000) NOREASSEMBLY))
  DSNAME('MWS.PQ33208.PTRACE4')
 PTRPT04I SNIFFER(ETHERNET) option selected

 OPTIONS((Both Bootp(67,68) Checksum(Summary) Cleanup(500) Datasize(0)
 DelayAck(200,200) Domain(53) EE(12000:12004) Finger(79) First Flags(All )
 Ftp(20,21) Gain(125,250) Gopher(70) Ike(500) Limit(999999999) Gmt
 Nat(4500) Ntp(123) Option Noreassembly Router(520) Rpc(111) Sasp(3860)
 Segment Smtp(25) Sniffer(4000,Ethernet) Snmp(161,162) Speed(10,10)
 Statistics(Detail) Telnet(23,80,) Tftp(69) Time(37) Tod Userexit() Www(80)
  ))
Sniffer Report
  3,385 records written to MWS.SNIFFER.ETH
639,789 bytes written
   2121 packets were abbreviated
   2024 is the maximum data size
   3375 packets were truncated from 1843 bytes
```

Following are descriptions for some areas of the example.

108 records written to SNIFFER
: This record count includes the two header records and one trailer record that were written to the SNIFFER data set.

46000 bytes written to SNIFFER
: The number of data bytes written to the SNIFFER data set. This should be close to the amount of data to be downloaded.

22 records were truncated to 1600 bytes
: Because the logical record length was 1,600 bytes, 22 records were truncated. This can be avoided by increasing the logical record length. The maximum logical record length is 32,763 or the size of physical disk blocks, whichever is smaller.

3.385 records written to MWS.SNIFFER.ETH
: This record count includes the two header records and one trailer record that were written to the SNIFFER data set.

639,789 bytes written
: The number of data bytes written to the SNIFFER data set. This should be the amount of data to be downloaded.

2121 packets were abbreviated
: The number of packets that were abbreviated when the trace data was collected.

2024 is the maximum data size
: The size of the largest record written to the SNIFFER data set.

3375 packets were truncated from 1843 bytes
: Because the logical record length was 2,048 bytes, 3375 records were truncated. This can be avoided by increasing the logical record length. The maximum logical record length is 32,763 or the size of physical disk blocks, whichever is smaller.

STATISTICS
Purpose

The records are counted by probe ID, device type, interface, interface address, job name, Asid, QOS, TCP port number, UDP port number, connection identifier, group identifier, type identifier, correlator, protocol summary, and session summary.

Format

The following command was used to obtain the example of this report.

```
CTRACE COMP(SYSTCPIS) SUB((TCPCS)) SHORT
OPTIONS((OPT STATISTICS(DETAIL)))
```

```
COMPONENT TRACE SHORT FORMAT
  SYSNAME(MVS118)
  COMP(SYSTCPIS)SUBNAME((TCPCS))
  OPTIONS((OPT STATISTICS(DETAIL)))
  DSNAME('IBMUSER.CTRACE1')

 OPTIONS((Both Bootp(67,68) Cleanup(500) DelayAck(200,200) Domain(53)
  Finger(79) Flags() Ftp(20,21) Gain(125,250) Gopher(70) Limit(999999999)
  Gmt Ntp(123) Option Noreassembly Router(520) Rpc(111) Segment Smtp(25)
  Snmp(161,162) Speed(10,10) Statistics(Detail) Telnet(23) Tftp(69) Time(37)
  Userexit() Www(80)
   ))

 **** 2002/11/20

 ================================================================================
1 SYSTCPIS Trace Statistics
  2,623 ctrace records processed
      0 segmented trace records read
      0 segmented trace records were lost
  2,623 trace records read
      0 records could not be validated
  2,623 records passed filtering
  2,623 packet trace records processed
      0 data trace records processed
 ================================================================================
2 Probe Report

   Total  Input    Data Output     Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Probe
    1526   1526      67144      0        0  4893 2002/11/20 17:56:00  7143 2002/11/20 18:09:17 03010021
       1      1         40      0        0  5652 2002/11/20 17:57:36  5652 2002/11/20 17:57:36 03010028
     859    859      34360      0        0  4553 2002/11/20 17:38:46  6376 2002/11/20 18:06:04 03020020
       6      6        724      0        0  4521 2002/11/20 17:38:32  5654 2002/11/20

.
.
.
    2623   2623     112084      0          0 Total

      9 Probe(s) found

3 Device Type Report

   Total  Input       Data Output   Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Device Type
     966    966      39300      0      0  4521 2002/11/20 17:38:32  6376 2002/11/20 18:06:04 1(LCS Ethernet)
     966    966      39300      0      0 Total

      1 Device Type(s) found

4 Interface Report

   Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Interface
     966    966      39300      0          0  4521 2002/11/20 17:38:32  6376 2002/11/20 18:06:04 ETH1
    1657   1657      72784      0          0  4522 2002/11/20 17:38:45  7143 2002/11/20 18:09:17 UNKNOWN
    2623   2623     112084      0          0 Total

      2 Interface(s) found

5 Interface Address Report

   Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Interface
     966    966      39300      0          0  4521 2002/11/20 17:38:32  6376 2002/11/20 18:06:04 ETH1
                                                    Addr: 9.42.104.38
    1557   1557      68384      0          0  4522 2002/11/20 17:38:45  7143 2002/11/20 18:09:17 UNKNOWN
                                                    Addr: 9.42.104.38
.
.
```

```
.
   2623   2623     112084      0          0 Total

    64 Interface Address(s) found
```

6 JobName Report

```
   Total  Input       Data Output   Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss JobName
    2610   2610     110984      0      0  4522 2002/11/20 17:38:45  7143 2002/11/20 18:09:17 FTPD1
       1      1         40      0      0  4587 2002/11/20 17:39:14  4587 2002/11/20 17:39:14 INETDCS1
       1      1        144      0      0  4591 2002/11/20 17:39:16  4591 2002/11/20 17:39:16 INETDCS3
       8      8        416      0      0  4521 2002/11/20 17:38:32  5892 2002/11/20 18:00:07 TCPCS
       1      1        123      0      0  4623 2002/11/20 17:40:48  4623 2002/11/20 17:40:48 TRMD
       2      2        377      0      0  5653 2002/11/20 17:57:37  5654 2002/11/20 17:57:37 USER17
    2623   2623     112084      0      0 Total
      6 JobName(s) found
```

7 Asid Report

```
   Total  Input       Data Output   Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Asid
    2623   2623     112084      0      0  4521 2002/11/20 17:38:32  7143 2002/11/20 18:09:17 01F7
    2623   2623     112084      0      0 Total
      1 Asid(s) found
```

8 Protocol Report

```
   Total  Input       Data Output    Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Protocol
      12     12        656      0       0  4521 2002/11/20 17:38:32  5892 2002/11/20 18:00:07 1(ICMP)
    2607   2607     110784      0       0  4522 2002/11/20 17:38:45  7143 2002/11/20 18:09:17 6(TCP)
       4      4        644      0       0  4591 2002/11/20 17:39:16  5654 2002/11/20 17:57:37 17(UDP)
    2623   2623     112084      0       0 Total   $

      3 Protocol(s) found
```

9 IP Address Report

```
   Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss
      11     11        484      0          0  6430 2002/11/20 18:09:02  7088 2002/11/20 18:09:16
                                                 Addr: 9.0.12.8
       1      1         40      0          0  4537 2002/11/20 17:38:45  4537 2002/11/20 17:38:45
                                                 Addr: 9.0.12.225
       1      1         56      0          0  5866 2002/11/20 18:00:06  5866 2002/11/20 18:00:06
                                                 Addr: 9.0.32.254
.
.
.
    5246   5246     224168      0          0 Total
    518 IP Address(s) found
```

10 Qos Report

```
   Total  Input       Data Output   Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Qos
       7      7        392      0      0  5830 2002/11/20 18:00:06  5892 2002/11/20 18:00:07 6(Internetwork)
       7      7        392      0      0 Total
      1 Qos(s) found
```

11 Tcp Port Report

```
   Total  Input       Data Output   Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Tcp Port
    2605   2605     110704      0      0  4522 2002/11/20 17:38:45  7143 2002/11/20 18:09:17 21(ftp)
       1      1         40      0      0  4743 2002/11/20 17:45:56  4743 2002/11/20 17:45:56 73(netrjs-3)
       1      1         40      0      0  5922 2002/11/20 18:00:10  5922 2002/11/20 18:00:10 74(netrjs-4)

.
.
.
   5214   5214     221568      0          0 Total

  1742 Tcp Port(s) found
```

12 Udp Port Report

```
   Total  Input       Data Output     Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Udp Port
       4      4        644      0       0  4591 2002/11/20 17:39:16  5654 2002/11/20 17:57:37 53(domain)
       1      1        144      0       0  4591 2002/11/20 17:39:16  4591 2002/11/20 17:39:16 1032()
       1      1        123      0       0  4623 2002/11/20 17:40:48  4623 2002/11/20 17:40:48 1033()
       1      1        144      0       0  5653 2002/11/20 17:57:37  5653 2002/11/20 17:57:37 1034()
       1      1        233      0       0  5654 2002/11/20 17:57:37  5654 2002/11/20 17:57:37 1035()
       8      8       1288      0       0 Total

      5 Udp Port(s) found
```

13 CID Report

```
   Total  Input       Data Output       Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss CID
     220    220       9200      0          0  4522 2002/11/20 17:38:45  5919 2002/11/20 18:00:07 00000020
       1      1         40      0          0  4553 2002/11/20 17:38:46  4553 2002/11/20 17:38:46 00000067
       1      1         40      0          0  4554 2002/11/20 17:38:46  4554 2002/11/20 17:38:46 00000096
.
.
.
    2615   2615     111668      0          0 Total
```

```
  2396 CID(s) found

14 Group Report
   Total  Input       Data Output    Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Group
    2423   2423     103508      0      0  4521 2002/11/20 17:38:32  7143 2002/11/20 18:09:17 3(SCAN)
     200    200       8576      0      0  5671 2002/11/20 17:59:32  5919 2002/11/20 18:00:07 4(ATTACK)
    2623   2623     112084      0      0 Total
      2 Group(s) found

15 Type Report
   Total  Input       Data Output    Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Type
    1527   1527      67184      0      0  4893 2002/11/20 17:56:00  7143 2002/11/20 18:09:17 0301(VSSCAN)
     859    859      34360      0      0  4553 2002/11/20 17:38:46  6376 2002/11/20 18:06:04 0302(PSSCAN)
      37     37       1964      0      0  4521 2002/11/20 17:38:32  5654 2002/11/20 17:57:37 0303(NORMSCAN)
     200    200       8576      0      0  5671 2002/11/20 17:59:32  5919 2002/11/20 18:00:07 0407(FLOOD)
    2623   2623     112084      0      0 Total
      4 Type(s) found

16 Correlator Report

   Total  Input       Data Output    Data First yyyy/mm/dd hh.mm.ss  Last yyyy/mm/dd hh.mm.ss Correlator
       4      4        644      0      0 4591 2002/11/20 17:39:16  5654 2002/11/20 17:57:37        2
       1      1         40      0      0  4521 2002/11/20 17:38:32  4521 2002/11/20 17:38:32       10
       1      1         40      0      0  4522 2002/11/20 17:38:45  4522 2002/11/20 17:38:45       11
.
.
.

    2623   2623     112084      0          0 Total

    467 Correlator(s) found

17 Protocol Summary Report
              Input                Output               Total
 Protocol Packets      Bytes  Packets      Bytes  Packets      Bytes
 Tcp         2607     110784        0          0     2607     110784
 Udp            4        644        0          0        4        644
 Icmp          12        656        0          0       12        656
 Other          0          0        0          0        0          0

18 Session Summary Report
  Input Output  First yyyy/mm/dd hh.mm.ss   Last yyyy/mm/dd hh.mm.ss Protocol
      1      0   5738 2002/11/20 17:59:34   5738 2002/11/20 17:59:34 TCP      Lcl:      9.4.81.167-27970
                                                                              Rmt:     9.42.104.38-21
      1      0   5710 2002/11/20 17:59:33   5710 2002/11/20 17:59:33 TCP      Lcl:     9.5.101.147-47426
                                                                              Rmt:     9.42.104.38-21
      1      0   5748 2002/11/20 17:59:34   5748 2002/11/20 17:59:34 TCP      Lcl:      9.6.159.21-30530
                                                                              Rmt:     9.42.104.38-21
.
.
.
   2618 session(s) found

   2623 records processed for this report
 Recording ended at   2002/11/20 18:09:17.543000
 Recording started at 2002/11/20 17:38:32.175560
 The duration was               00:30:45.367440
      1 records with ABBREV=200
   2622 records with FULL=144
    233 is the maximum packet data length
 655360 bytes of storage used to create this report
   7841 requests for 652704 bytes of storage were made
```

The following describes numbered areas of the example.

1

The standard statistics shown with all executions of the SYSTCPIS packet trace formatter.

2,623 ctrace records processed
The total number of CTRACE records given to the SYSTCPIS packet trace formatted.

0 segmented trace records read
The total number of packets that spanned multiple CTRACE records.

0 segmented trace records were lost
The total number of packets records that could not be put back together.

2,623 trace records read
The total number of complete trace records.

0 records could not be validated
The number of incomplete CTRACE records that could not be used.

2,623 records passed filtering
The number of records that were successfully formatted.

2,623 packet trace records processed
The number of records that were packet trace records.

0 data trace records processed
The number of records that were data trace records.

2
Probe report, which is the total by ProbeID.

3
Device type report, which is the totals by device type.

4
Interface report, which is the totals by interface.

5
Interface address report, which is the totals interface address.

6
Jobname report, which is the totals by jobname.

7
ASID report, which is the totals address space identifier.

8
Protocol report, which is the totals by protocol.

9
IP address report, which is the totals by IP address. Both the destination and source IP addresses are counted, except when they are the same in a record.

10
QOS report, which is the totals by QOS.

11
TCP port report, which is the totals by TCP port number. Both the destination and source port numbers are counted, except when they are the same in a record.

12
UDP port report, which is the totals by UPD port number. Both the destination and source port numbers are counted, except when they are the same in a record.

13
CID report, which is the totals by connection identifier.

14
Group report, which is the totals by group, first byte PROBEID.

15
Type report, which is the totals by type, first two bytes of PROBEID.

16
Correlator report, which is the totals by correlator.

17
Protocol summary report, which is the summary based on protocol.

18
Session summary report, which is the summary based on session.

STREAM
Purpose

There are times when messages span multiple packets. TELNET and DNS are examples. The STREAM report (with the DUMP or FORMAT keywords) captures the entire stream of data.

Format

The following command was used to obtain the example of this report.

```
CTRACE COMP(SYSTCPIS) SUB((TCPCS)) SHORT DSN('IBMUSER.CTRACE1') OPTIONS((OPT STREAM DUMP ASCII))

COMPONENT TRACE SHORT FORMAT
 SYSNAME(MVS118)
 COMP(SYSTCPIS)SUBNAME((TCPCS))
 OPTIONS((OPT STREAM DUMP ASCII))
 DSNAME('IBMUSER.CTRACE1')

OPTIONS((Ascii Bootp(67,68) Cleanup(500) DelayAck(200,200) Domain(53)
 Dump(65535) Finger(79) Flags() Ftp(20,21) Gain(125,250) Gopher(70)
 Limit(999999999) Gmt Ntp(123) Option Noreassembly Router(520) Rpc(111)
 Segment Smtp(25) Snmp(161,162) Speed(10,10) Streams(131072,Summary)
 Telnet(23) Tftp(69) Time(37) Userexit() Www(80)
  ))

**** 2002/11/20
RcdNr Sysname  Mnemonic Entry Id   Time Stamp    Description
----- -------- -------- -------- --------------- ------------------------------

================================================================================
1 Streams Report
  2618 Streams found
611952 bytes of storage for the session report was allocated
348160 bytes of storage for buffers was allocated
 --------------------------------------------------------------------------------
.
.
.
 --------------------------------------------------------------------------------
2 Session: 9.32.74.253-0 9.42.104.38-0 ICMP
  From: 2002/11/20 18:00:06.827658 to: 2002/11/20 18:00:07.149355
       2 packets found
  Stream buffer at 16743000 for 20480 bytes. 56 bytes were used
       2 packets moved for 56 bytes
 I - Inbound packet
 O - Outbound packet

3 D Rcd #            Time           Delta      Seq #   Position Length    End_Pos
 I  5870 18:00:06.827658 00:00:00.000000          0          0     28         28
 000000 45000028 1F9B0000 0106EC56 092A6826  09A032EF 0015E644 7F6CBB58 |E..(.......V.*h&..2....D.l.X    |
 I  5892 18:00:07.149355 00:00:00.321697         28         28     28         56
 000000                                                  4500002C |                              E..,|
 000020 1FDC0000 0106EC11 092A6826 09A032EF  00154446 809241F5              |.........*h&..2...DF..A.        |


.
.
.
 ================================================================================
 SYSTCPIS Trace Statistics
  2,623 ctrace records processed
      0 segmented trace records read
      0 segmented trace records were lost
  2,623 trace records read
      0 records could not be validated
  2,623 records passed filtering
  2,623 packet trace records processed
      0 data trace records processed
```

OSAENTA trace (SYSTCPOT)

TCP/IP Services component trace is also available for use with the OSA-Express Network Traffic Analyzer (OSAENTA) trace facility. The OSAENTA trace is a diagnostic method for obtaining frames flowing to and from an OSA adapter. You can use the OSAENTA statement to copy frames as they enter or leave an OSA

adapter for an attached host. The host can be an LPAR with z/OS, VM, or Linux. You can then examine the contents of the copied frames. To be traced, the frame must meet all the conditions specified on the OSAENTA statement or the OSAENTA command.

The OSAENTA trace process

Trace data is collected as frames enter or leave an OSA adapter for a connected host. The actual collection occurs within the device drivers of OSA cards, capturing the data at the point where it has just been received from or sent to the network.

Frames that are captured have extra information added to them before they are stored. This extra information, such as timestamps, is used during the packet formatting. The captured data reflects exactly what the network sees. For example, the trace contains the constituent packets of a fragmented packet exactly as they are received or sent.

The selection criteria for choosing packets to trace are specified through the OSAENTA statement or OSAENTA command. See z/OS Communications Server: IP Configuration Reference for more information about the OSAENTA statement and see z/OS Communications Server: IP System Administrator's Commands for more information about the OSAENTA command.

The OSAENTA trace can have performance implications if you do not specify sufficient trace filters before enabling the trace. OSAENTA can reduce the amount of traffic the OSA-Express feature can process and the amount of traffic that can be accelerated through that OSA-Express. Also, host processing to collect the OSAENTA trace records can increase host CPU consumption. Specify sufficient filters to limit the amount of traffic that is traced to only what is necessary for problem diagnosis.

Figure 25 illustrates the overall control and data flow in the OSAENTA tracing facility.

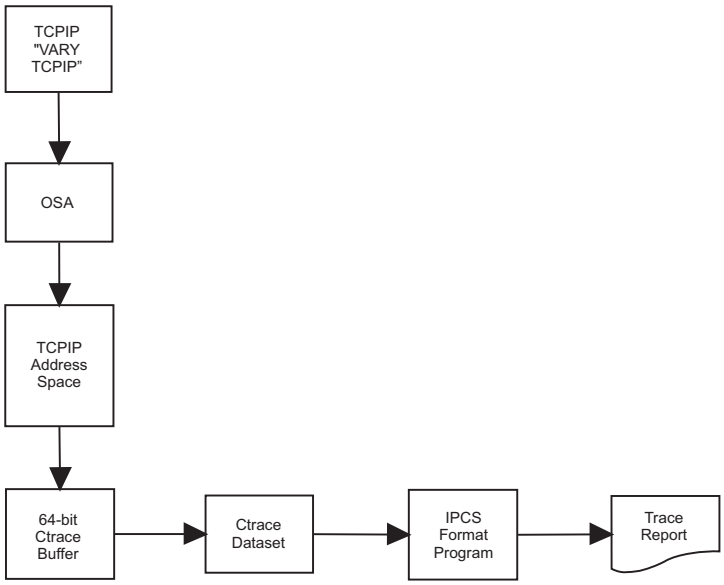

Figure 25. Control and data flow in the OSAENTA tracing facility

Starting OSAENTA trace

You can start an OSAENTA trace in one of the following ways:

- Using the V TCPIP,,OSAENTA command

  ```
  V TCPIP,tcpprocname,OSAENTA,ON,PORTNAME=OSA4,IPADDR=9.1.27.2
  ```

- Using the OSAENTA statement in TCPIP.PROFILE

  ```
  OSAENTA ON PORTNAME=OSA4 IPADDR=9.1.27.2
  ```

Security Rule: To use any VARY command, the user must be authorized in RACF. The OPERCMDS RACF profile for each user must have access for a resource of the form MVS.VARY.TCPIP.OSAENTA.

Traces are placed in an internal buffer, which can then be written out using a CTRACE external writer. The MVS TRACE command must also be issued for component SYSTCPOT to activate the OSAENTA trace.

After starting OSAENTA trace, you can display the status using the Netstat command, as shown in the following example:

```
D TCPIP,TCPCS,NETSTAT,DEV
  DEVNAME: OSA4              DEVTYPE: MPCIPA
    DEVSTATUS: READY
    LNKNAME: LOSAFE            LNKTYPE: IPAQENET   LNKSTATUS: READY
.
.
.
OSA-Express Network Traffic Analyzer Information:
  OSA PortName: OSA4             OSA DevStatus:     Ready
    OSA IntfName: EZANTAQDIO4101  OSA IntfStatus:    Ready
    OSA Speed:    1000            OSA Authorization: Logical Partition
    OSAENTA Cumulative Trace Statistics:
      DataMegs:   0                     Frames:          8
      DataBytes:  760                   FramesDiscarded: 4
      FramesLost: 0
    OSAENTA Active Trace Statistics:
      DataMegs:   0                     Frames:          8
      DataBytes:  760                   FramesDiscarded: 4
      FramesLost: 0                     TimeActive:      8
    OSAENTA Trace Settings:           Status: On
      DataMegsLimit: 1024               FramesLimit:     2147483647
      Abbrev:        224                TimeLimit:       10080
      Discard:       ALL
    OSAENTA Trace Filters:            Nofilter: ALL
      DeviceID: *
      Mac:      *
      VLANid:   *
      ETHType:  *
      IPAddr:   *
      Protocol: *
      PortNum:  *
```

If you are a TSO user, use the NETSTAT DEVlinks command.

Modifying options with VARY commands

After starting an OSAENTA trace, you can change the trace using the VARY command. For example, if you want to change the trace to abbreviate the data being traced, use the following command:

V TCPIP,*tcpipproc*,OSAENTA,ON,ABBREV=480

You can display the results of the VARY command using Netstat:

```
netstat -p TCPCS -d
  DEVNAME: OSAQDIO4          DEVTYPE: MPCIPA
    DEVSTATUS: READY
    LNKNAME: LOSAFE            LNKTYPE: IPAQENET   LNKSTATUS: READY
.
.
.
OSA-Express Network Traffic Analyzer Information:
  OSA PortName: QDIO4101          OSA DevStatus:     Ready
    OSA IntfName: EZANTAQDIO4101  OSA IntfStatus:    Ready
    OSA Speed:    1000            OSA Authorization: Logical Partition
    OSAENTA Cumulative Trace Statistics:
      DataMegs:   0                     Frames:          8
      DataBytes:  760                   FramesDiscarded: 4
      FramesLost: 0
    OSAENTA Active Trace Statistics:
      DataMegs:   0                     Frames:          8
      DataBytes:  760                   FramesDiscarded: 4
      FramesLost: 0                     TimeActive:      8
    OSAENTA Trace Settings:           Status: On
      DataMegsLimit: 1024               FramesLimit:     2147483647
      Abbrev:        480                TimeLimit:       10080
      Discard:       ALL
    OSAENTA Trace Filters:            Nofilter: ALL
      DeviceID: *
      Mac:      *
      VLANid:   *
      ETHType:  *
      IPAddr:   *
      Protocol: *
      PortNum:  *
```

If you are a TSO user, use the NETSTAT DEVlinks option.

You can use the VARY TCPIP,*tcpproc*,OBEYFILE command to make temporary dynamic changes to system operation and network configuration without stopping and restarting the TCP/IP address space. For example, if you started the address space TCPIPA and created a sequential data set USER99.TCPIP.OBEYFIL2 containing OSAENTA statements, issue the following command:

```
VARY TCPIP,TCPIPA,CMD=OBEYFILE,DSN=USER99.TCPIP.OBEYFIL2
```

The VARY TCPIP,,OSAENTA command is cumulative. You can trace all packets for specified IP addresses by entering multiple OSAENTA commands. In the following example, the five commands disable the current trace, clear any previous trace filters, trace all the frames received and all the frames sent for the specified IP addresses, and activate the OSAENTA trace facility.

```
VARY TCPIP,,OSAENTA,OFF,PORTNAME=OSA4
VARY TCPIP,,OSAENTA,CLEAR,PORTNAME=OSA4,ABBREV=200,FRAMES=8000
VARY TCPIP,,OSAENTA,PORTNAME=OSA4,IPADDR=10.27.142.44
VARY TCPIP,,OSAENTA,PORTNAME=OSA4,IPADDR=10.27.142.45
VARY TCPIP,,OSAENTA,ON,PORTNAME=OSA4
```

Formatting OSA traces using IPCS

The IPCS CTRACE command parameters are described in "Formatting component traces" on page 53. The following notes apply to the IPCS CTRACE parameters with regard to the OSAENTA trace formatter:

JOBLIST, JOBNAME
: The LINKNAME and JOBNAME keywords in the OPTIONS string can also be used to select records.

TALLY
Equivalent to the STATISTICS(DETAIL) option.

START, STOP
The time is set when the record was moved to the trace buffer, not when the OSA card recorded the data.

LIMIT
See the RECORDS keyword in the OPTIONS string.

USEREXIT
The packet trace formatter calls the CTRACE USEREXIT before testing the records with the filtering criteria. If it returns a nonzero return code, then the record is skipped. The USEREXIT can also be used in the OPTIONS string. It is called after the record has met all the filtering criteria in the OPTIONS string.

COMP
Must be SYSTCPOT.

SUB
The SUB must name the TCP/IP procedure that created the CTRACE records when the input is a dump data set.

EXCEPTION
Since there are no EXCEPTION records for OSAENTA trace, the EXCEPTION keyword must not be specified.

ENTIDLIST
The following values are the valid values for OSAENTA trace:

```
7  Link Frame trace records
```

The CTRACE OPTIONS string provides a means of entering additional keywords for record selection and formatting OSA traces (COMP=SYSTCPOT). See "Syntax" on page 55 for the complete syntax of CTRACE.

The same program is used to format OSA traces as well as packet traces. See "OPTIONS syntax" on page 97 for the values specified for the OPTIONS keyword.

Network security services (NSS) server trace (SYSTCPNS)

TCP/IP Services component trace is also available for use with the network security services server. See "TCP/IP services component trace for the network security services (NSS) server" on page 377.

Defense Manager daemon (DMD) trace (SYSTCPDM)

TCP/IP Services component trace is also available for use with the Defense Manager daemon (DMD). See "TCP/IP services component trace for the Defense Manager daemon" on page 728.

OMPROUTE trace (SYSTCPRT)

TCP/IP Services component trace is also available for use with the OMPROUTE application. See "TCP/IP services component trace for OMPROUTE" on page 786.

RESOLVER trace (SYSTCPRE)

TCP/IP Services component trace is also available for use with the RESOLVER application. See Chapter 39, "Diagnosing resolver problems," on page 869.

Configuration profile trace

You can use the ITRACE statement in the PROFILE.TCPIP data set to activate TCP/IP run-time tracing for configuration, the TCP/IP SNMP subagent, commands, and the autolog subtask. ITRACE should only be set at the direction of an IBM Support Center representative.

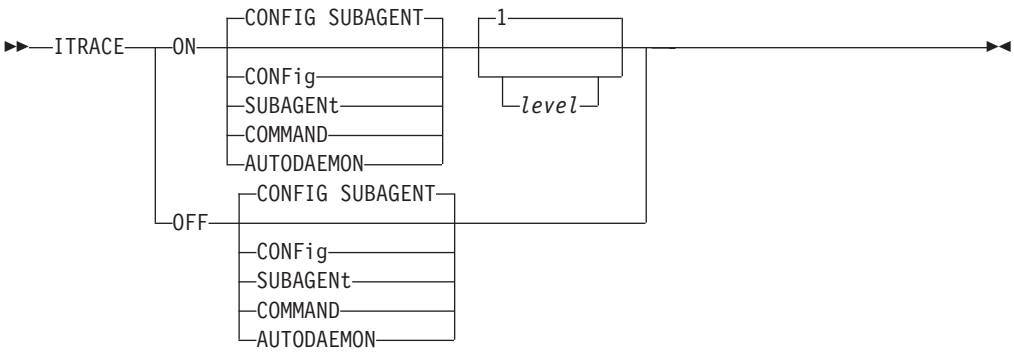

Following are descriptions of the ITRACE parameters:

ON Select ON to establish run-time tracing. ITRACE ON commands are cumulative until an ITRACE OFF is issued.

OFF
Select OFF to terminate run-time tracing.

CONFig
Turn internal trace for configuration ON or OFF.

SUBAgent
Turn internal trace for TCP/IP SNMP subagent ON or OFF.

COMMAND
Turn internal trace for command ON or OFF.

AUTODAEMON
Turn internal trace for the autolog subtask ON or OFF.

level
Indicates the tracing level to be established. Levels are as follows:

Levels for CONFIG

1 ITRACE for all of config

2 General level of tracing for all of config

3 Tracing for configuration set commands

4 Tracing for configuration get commands

5 Tracing for syslog calls issued by config

100 Tracing for the parser

200 Tracing for scanner

300 Tracing for mainloop

400 Tracing for commands

Levels for SUBAGENT

1 General subagent tracing

2 General subagent tracing, plus DPI traces

3 General subagent tracing, plus extended storage dump traces

4 All trace levels

Level for COMMAND

1 ITRACE for all commands

Following is an example illustrating how to use the ITRACE command:

```
ITRACE ON CONFIG 3
ITRACE OFF SUBAGENT
```

Trace output is sent to the following locations:

- Subagent trace output is directed to the syslog daemon. This daemon is configured by the /etc/syslog.conf file and must be active.
- AUTOLOG trace output goes to ALGPRINT.
- Trace output for other components goes to SYSPRINT.

Chapter 6. IPCS subcommands for TCP/IP

Use the IPCS subcommands for TCP/IP to format data from IPCS system dumps. This topic describes the subcommands (including description, syntax, parameters, and sample output), installation, entering, and execution, and includes the following subtopics:

- "TCPIPCS command" on page 176
- "TCPIPCS subcommands" on page 179
- "ERRNO command" on page 276
- "IPHDR command" on page 280
- "RESOLVER command" on page 282
- "SETPRINT command" on page 303
- "SKMSG command" on page 303
- "TCPHDR command" on page 305
- "TOD command" on page 306
- "UDPHDR command" on page 307
- "Installing TCP/IP IPCS subcommands by using the panel interface" on page 309
- "Entering TCP/IP IPCS subcommands" on page 309

Types of subcommands

There are two types of subcommands. These are described as follows:

- Many of the TCP/IP subcommands work on a specific stack or Telnet instance. These subcommands are grouped under the TCPIPCS subcommand to share the TCP (to select the stack or Telnet) and TITLE options. A subset of these commands are available for work with an instance of Telnet. If available, "Available for Telnet" appears at the end of the description in Table 15.
- The remaining TCP/IP IPCS subcommands do not require a TCP/IP stack, and they are not under the TCPIPCS subcommand.

Restriction: The TCP/IP IPCS commands are not supported for IPCS "active."

Table 15 lists all the IPCS subcommands. The TCPIPCS commands are shown first, followed by the general commands.

Table 15. TCP/IP IPCS commands

Command	Description	Alias	See
TCPIPCS ALL	Equivalent to TCPIPCS STATE TSEB TSDB TSDX DUAF CONFIG ROUTE SOCKET STREAM RAW TCB UDP LOCK TIMER STORAGE		NA
TCPIPCS API	Display control blocks for Sockets Extended Assembler Macro and Pascal APIs		"TCPIPCS API" on page 179

© Copyright IBM Corp. 2000, 2015

Table 15. TCP/IP IPCS commands (continued)

Command	Description	Alias	See
TCPIPCS CONFIG	Display device configuration information	TCPIPCS CNFG TCPIPCS CONF	"TCPIPCS CONFIG" on page 181
TCPIPCS CONNECTION	Display active or all connections	TCPIPCS CONN	"TCPIPCS CONNECTION" on page 182
TCPIPCS COUNTERS	Display information about TCP/IP internal execution statistics		"TCPIPCS COUNTERS" on page 184
TCPIPCS DETAIL	Equivalent to TCPIPCS TSEB TSDB TSDX DUAF Available for Telnet.	TCPIPCS CBS	NA
TCPIPCS DU	Equivalent to TCPIPCS DUAF DUCB Available for Telnet.		NA
TCPIPCS DUAF	Summarize DUCBs Available for Telnet.	TCPIPCS DUCBS	"TCPIPCS DUAF" on page 186
TCPIPCS DUCB	Find and format DUCBs Available for Telnet.		"TCPIPCS DUCB" on page 189
TCPIPCS FRCA	Display state information about FRCA connections and objects		"TCPIPCS FRCA" on page 192
TCPIPCS HASH	Display TCP/IP data stored in hash tables		"TCPIPCS HASH" on page 194
TCPIPCS HEADER	Display dump Header info	TCPIPCS HDR	"TCPIPCS HEADER" on page 198
TCPIPCS HELP	Display syntax help for TCPIPCS command	TCPIPCS ?	"TCPIPCS HELP" on page 200
TCPIPCS IPSEC	Display information about IP security filters and tunnels		"TCPIPCS IPSEC" on page 201
TCPIPCS LOCK	Display locks Available for Telnet.	TCPIPCS LOCKSUM	"TCPIPCS LOCK" on page 204
TCPIPCS MAP	Display storage map		"TCPIPCS MAP" on page 205
TCPIPCS MTABLE	Display module table		"TCPIPCS MTABLE" on page 208
TCPIPCS POLICY	Display service policy data		"TCPIPCS POLICY" on page 209
TCPIPCS PROFILE	Display TCP/IP configuration data in the format of a profile dataset Available for Telnet.	TCPIPCS PROF	"TCPIPCS PROFILE" on page 211

Table 15. TCP/IP IPCS commands (continued)

Command	Description	Alias	See
TCPIPCS PROTOCOL	Invokes RAW, TCB, UDP		"TCPIPCS PROTOCOL" on page 215
TCPIPCS RAW	Display RAW control blocks	TCPIPCS MRCB TCPIPCS RAWSUM TCPIPCS RCB	"TCPIPCS RAW" on page 218
TCPIPCS ROUTE	Display routing information	TCPIPCS RTE	"TCPIPCS ROUTE" on page 220
TCPIPCS SOCKET	Display socket information	TCPIPCS SCB TCPIPCS SOCKSUM	"TCPIPCS SOCKET" on page 223
TCPIPCS STATE	Display general stack information	TCPIPCS	"TCPIPCS STATE" on page 225
TCPIPCS STORAGE	Display TCP/IP storage usage	TCPIPCS STOR	"TCPIPCS STORAGE" on page 246
TCPIPCS STREAM	Display streams information	TCPIPCS SKSH TCPIPCS STREAMS	"TCPIPCS STREAM" on page 248
TCPIPCS SUMMARY	Equivalent to TCPIPCS DUAF CONFIG SOCKET		NA
TCPIPCS TCB	Display TCP protocol control blocks	TCPIPCS MTCB TCPIPCS TCBSUM	"TCPIPCS TCB" on page 250
TCPIPCS TELNET	Display Telnet information Available for Telnet.		"TCPIPCS TELNET" on page 252
TCPIPCS TIMER	Display information about timers Available for Telnet.	TCPIPCS TIMESUM	"TCPIPCS TIMER" on page 254
TCPIPCS TRACE	Display TCP/IP CTRACE, real-time TCP/IP network monitoring NMI (TMI), and real-time application-controlled TCP/IP trace NMI information Available for Telnet.	TCPIPCS TCA	Table 3 on page 8
TCPIPCS TREE	Display information about data stored in Patricia trees Available for Telnet.	TCPIPCS TREESUM	"TCPIPCS TREE" on page 260
TCPIPCS TSDB	Format TSDB Available for Telnet.		"TCPIPCS TSDB" on page 263
TCPIPCS TSDX	Format TSDX Available for Telnet.		"TCPIPCS TSDX" on page 265

Table 15. TCP/IP IPCS commands (continued)

Command	Description	Alias	See
TCPIPCS TSEB	Format TSEB Available for Telnet.		"TCPIPCS TSEB" on page 266
TCPIPCS TTLS	Display state information about AT-TLS connections and groups		"TCPIPCS TTLS" on page 267
TCPIPCS UDP	Display UDP control blocks	TCPIPCS MUCB TCPIPCS UCB TCPIPCS UDPSUM	"TCPIPCS UDP" on page 270
TCPIPCS VMCF	Display information about VMCF and IUCV users		"TCPIPCS VMCF" on page 273
TCPIPCS XCF	Display information about XCF links and dynamic VIPA		"TCPIPCS XCF" on page 274
ERRNO	Interpret error numbers Available for Telnet.		"ERRNO command" on page 276
ICMPHDR	Format an ICMP header		"ICMPHDR command" on page 278
IPHDR	Format an IP header		"IPHDR command" on page 280
RESOLVER	Format and summarize resolver control blocks and cache information		"RESOLVER command" on page 282
SETPRINT	Set destination so the IPCS subcommand output is sent to a user ID or the printer Available for Telnet.		"SETPRINT command" on page 303
SKMSG	Format a stream message Available for Telnet.		"SKMSG command" on page 303
TCPHDR	Format a TCP header		"TCPHDR command" on page 305
TOD	Convert an S/390® 64-bit time-of-day timestamp to a readable date and time Available for Telnet.		"TOD command" on page 306
UDPHDR	Format UDP header		"UDPHDR command" on page 307

TCPIPCS command

This topic describes the TCPIPCS command.

Syntax

The command syntax for all TCPIPCS subcommands includes an option to specify the TCP stack and to specify whether the title is displayed.

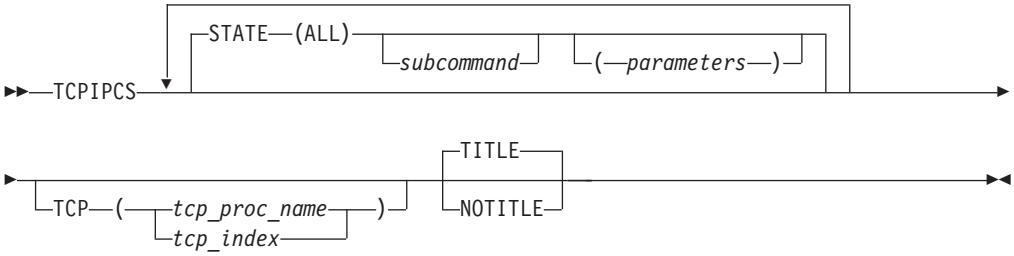

Parameters

The parameters for the TCPIPCS command are described below.

subcommand
Default is STATE.

parameters
Each subcommand has its own parameters.

- If a command has variable parameters, they can be omitted, specified as a single variable, or specified as a list. If no variable parameters are specified, an asterisk must be used as a placeholder if any keyword parameters are specified. If two or more variable parameters are specified, they must be enclosed in parentheses.
- To distinguish between the variable parameters, a parameter is assumed to be one of the following:
  - An index or small number if it is four digits or less, begins with zero to nine, and contains only hexadecimal digits (0–9, a–f, A–F). If a command accepts multiple indices or small numbers, both are compared to the values and the first matching field is used.
  - An address if it is more than four digits, begins with zero to nine, and contains only hexadecimal digits. For example, for the TCPIPCS DUAF command, both the DUCB and ASCB addresses of each DUCB are compared to the address parameter, and the first matching field is used to select the DUCB to display.
  - An IPCS symbol name can also be specified for an address.
  - Otherwise, the parameter is assumed to be a character string variable (such as TCP/IP procedure or job name, user ID, and command name).
- Keyword parameters can be in any order.
- If there are both keyword and variable parameters, all variable parameters must precede the keywords.

TCP
Specifies which TCP/IP stack or Telnet instance. When issuing commands for Telnet, the Telnet procedure name must be specified in the tcp_proc_name variable. The stack can be specified directly or indirectly. A stack can be specified directly by coding the TCP parameter with either *tcp_proc_name* or *tcp_index*. If no stack is specified directly, the output is reported for the stack with the lowest index matching the release of the TCPIPCS command. After a

particular stack is specified (whether specified directly or indirectly), that stack becomes the default. The stack index is saved as a symbol and is used as the default in future invocations of the TCPIPCS command. An alias for the TCP option is PROC.

Note: All eight stack indices are available when TCP/IP or Telnet starts, so any stack index can be selected. The existence of an index does not necessarily mean this stack can be formatted. If the stack was not included in the dump, then most of the information about a stack cannot be formatted. Most TCP/IP control blocks are in the private TCP address space. All Telnet control blocks are in the private Telnet address space.

The fact that an index exists does not necessarily mean this stack index has ever been used. If you specify a stack index that has not been used, the version and release fields for this stack are zero, so you receive a message indicating the stack is not the same version and release as the TCPIPCS command:

tcp_proc_name
TCP/IP procedure name or the Telnet procedure name (when the TN3270E Server is running in its own address space).

tcp_index
TCP/IP stack index (1-8) or Telnet index (9-16).

TITLE
The title contains information about the dump and about the TCPIPCS command. By default, the title information is displayed.

The title contains the following information.

- TCPIPCS command input parameters.
- Dump data set name.
- Dump title.
- TSAB address.
- Table listing all TCP/IP stacks used in the dump and their
  - TSEB address
  - Stack index
  - Procedure name
  - Stack version
  - TSDB address
  - TSDX address
  - ASID
  - Trace option bits
  - Stack status
- Count of the number of TCP/IP stacks defined (used).
- Count of the number of active TCP/IP stacks found.
- Count of the number of active TCP/IP stacks matching the TCPIPCS command version and release.
- Procedure name and index of the stack being reported.

NOTITLE
Suppress the title lines. This is useful when you are processing many commands on the same dump and do not want to see the title information repeated.

Rule: If you specify multiple keywords from the set {TITLE, NOTITLE}, only the last one is used.

Symbols defined

TCPIPCS defines the following IPCS symbols:

TSEBPTR
The address of the first TSEB control block.

TSEB*n*
The address of the TSEB control block corresponding to the stack index *n*.

TCPIPCS subcommands

This topic describes the TCPIPCS subcommands.

TCPIPCS API

Use this subcommand to display information about the connections in the Sockets Extended Assembler Macro Application Programming Interface (Macro API) and the Pascal API.

Note: The Macro API is the base for the CALL Instruction API, the CICS C API, and the CICS EZACICAL API. See the z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference for more information about the native TCP/IP APIs.
Some API control blocks are in the application address space, which might not be available in the dump. If the application address space is available, the API control blocks are formatted.

Syntax

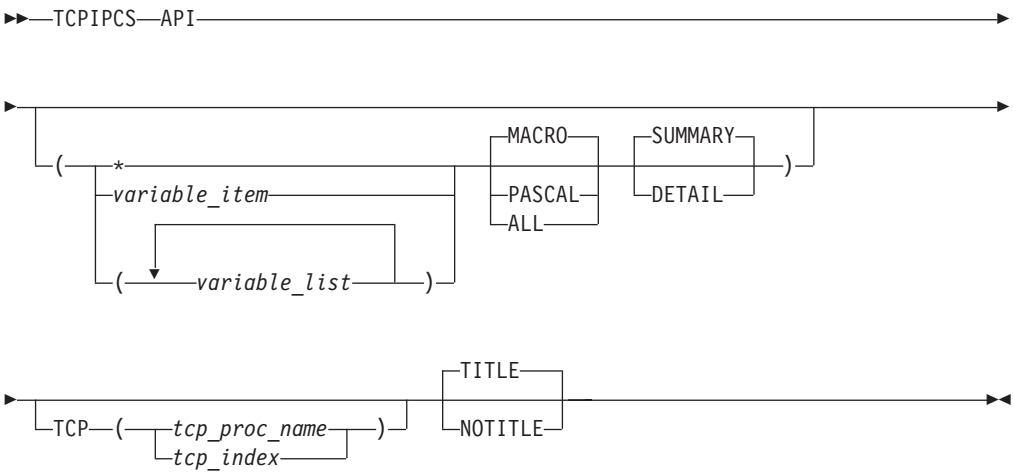

Parameters

If no parameters are specified, only information about the Macro API is summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 – 32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

jobname
Displays only the API control blocks for this job name. The job name can be a TCP/IP application name or a stack name. Must contain from 1-8 characters.

ASCB_address
Displays the API control blocks with this address space control block (ASCB) address. An IPCS symbol name can be specified for the address. The address is specified as 1-8 hexadecimal digits. If an address begins with digit A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

ASID_number
Displays the API control blocks with this address space identifier (ASID). The ASID is a hexadecimal number containing 1 - 4 digits.

In addition to the variable parameters, you can specify the following keyword parameters:

MACRO
Displays only information for Macro APIs. MACRO is the default.

PASCAL
Displays only information for Pascal APIs.

ALL
Displays information for both APIs.

SUMMARY
Displays the addresses of the control blocks and other data in tables. SUMMARY is the default.

DETAIL
Also displays the contents of the control blocks in addition to the SUMMARY display.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {MACRO, PASCAL, ALL}, only the last one is used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS API subcommand

The following example is sample output of the TCPIPCS API subcommand.

```
The contents of the SDST control blocks are formatted by the TCPIPCS API
subcommand if the DETAIL option is coded on the command (SUMMARY is the
default and only the address of the SDST will be displayed in this case).


R14 Output:

        -- Array elements --
   ::
```

```
+00B2  SDST_LOCAL_IPADDRLEN.  00
+00B3  SDST_REMOTE_IPADDRLEN. 00
+00B4  SDST_LOCAL_IPADDR.     00000000  00000000  00000000  00000000
+00C4  SDST_REMOTE_IPADDR.    00000000  00000000  00000000  00000000
::
     -- End of array   --
```

TCPIPCS CONFIG

Use this subcommand to display each device interface, physical interface, and logical interface. The configuration summary table shows each logical interface with the name of its associated device and link.

Syntax

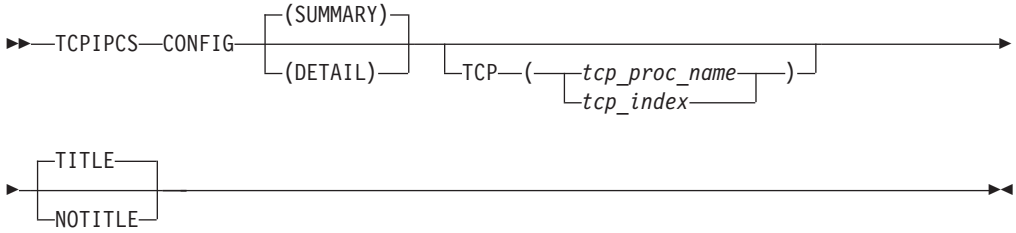

Parameters

SUMMARY
Displays each device, physical interface, and logical interface, and summarizes them all in one cross-reference table. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows the interface cross-reference reports.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS CONFIG subcommand

The following example is sample output of the TCPIPCS CONFIG subcommand.

```
TCPIPCS CONFIG
Dataset: IPCS.R450697.V6TCBD1
Title:   TCPCS2 CLIENT SIDE


The address of the TSAB is: 09DBE1A0

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

09DBE1E0  1 TCPCS     V1R5    096C4000 096C40C8 0033 10841004  Active
09DBE260  2 TCPCS2    V1R5    096C9000 096C90C8 0034 10841004  Active

   2 defined TCP/IP(s) were found
   2 active  TCP/IP(s) were found

   2 TCP/IP(s) for CS     V1R5  found

================================================================================

Analysis of Tcp/Ip for TCPCS2.  Index: 2
```

```
Configuration control block summary

IPMAIN found at 095A83D0

IPMAIN6 found at 096CE470

Dif@      DeviceName       Next     Prev     DevR DevW Protocol
7F6AA408  LOOPBACK         7F1ED408 00000000 **** **** LOOPBACK
7F1ED408  OSAQDIO3         00000000 7F6AA408 **** **** MPCIPA

IPv4 Pif@ LinkName         Next     Prev     DeviceName       Protocol        Dif@     Lif@
7F679468  LOOPBACK         7F503B88 00000000 LOOPBACK         LOOPBACK        7F6AA408 7F6792E8
7F503B88  OSAQDIOL         00000000 7F679468 OSAQDIO3         IPAQENET        7F1ED408 7F1ED008

IPv6 Pif@ IntfName         Next     Prev     DeviceName       Protocol        Dif@     Lif@
7F503488  LOOPBACK6        7F503F08 00000000 LOOPBACK         LOOPBACK6       7F6AA408 7F3FDCE8
7F503F08  OSAQDI26         00000000 7F503488 OSAQDIO3         IPAQENET6       7F1ED408 7F6BF028

IPv4 Lif@ LinkName         Next     Prev     Pif@     IpAddr
7F1ED008  OSAQDIOL         7F6792E8 00000000 7F503B88 9.67.115.82
7F6792E8  LOOPBACK         00000000 7F1ED008 7F679468 127.0.0.1

IPv6 Lif@ IntfName         Next     Prev     Pif@     IpAddr
7F3083C8  OSAQDI26         7F3FDCE8 00000000 7F503F08 FEC9:C2D4:1::9:67:115:82
7F3FDCE8  LOOPBACK6        7F6BF028 7F3083C8 7F503488 ::1
7F6BF028  OSAQDI26         00000000 7F3FDCE8 7F503F08 FE80::2:559A:3F5F:1

Configuration Summary

IPv4 Lif@ LinkName         DeviceName       DevR DevW IpAddr
7F1ED008  OSAQDIOL         OSAQDIO3         **** **** 9.67.115.82
7F6792E8  LOOPBACK         LOOPBACK         **** **** 127.0.0.1

IPv6 Lif@ LinkName         DeviceName       DevR DevW IpAddr
7F3083C8  OSAQDI26         OSAQDIO3         **** **** FEC9:C2D4:1::9:67:115:82
7F3FDCE8  LOOPBACK6        LOOPBACK         **** **** ::1
7F6BF028  OSAQDI26         OSAQDIO3         **** **** FE80::2:559A:3F5F:1

Analysis of Tcp/Ip for TCPCS2 completed
```

TCPIPCS CONNECTION

Use this subcommand to display information about TCP, UDP, and raw connections. The information includes the following:

- User ID
- Connection ID
- Local IP address
- Foreign IP address
- Connection state (for TCP connections only)
- Protocol name (for raw connections only)

Syntax

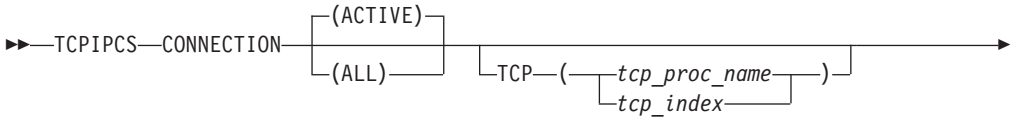

Parameters

ACTIVE
Display only active connections. This is the default.

Tip: The number of connections reported for each protocol includes both inactive and active connections; therefore, the total might be higher than the number of displayed (active) connections.

ALL
Display all connections, regardless of state.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {ACTIVE, ALL}, only the last one is used.

Sample output of the TCPIPCS CONNECTION subcommand

The following is a sample output of the TCPIPCS CONNECTION subcommand. In this sample, the default option is ACTIVE, so only active connections are shown. There are six active TCP connections, four active UDP connections, and three active RAW connections.

```
TCPIPCS CONNECTION
Dataset: IPCS.R8A0723.RASDUMP
Title:   EZRPE005
The address of the TSAB is: 098221F0
Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status
09822230  1 TCPCS     V1R5    08E85000 08E850C8 001E 9FFF7E7F  Active
098222B0  2 TCPCS2    V1R5    08937000 089370C8 01F6 9FFF7E7F  Active
   2 defined TCP/IP(s) were found
   2 active  TCP/IP(s) were found
   2 TCP/IP(s) for CS     V1R5  found
================================================================================
Analysis of Tcp/Ip for TCPCS.  Index: 1

TCP IPv4 Connections:
Userid   Conn       State
TCPCS    00000012   Listening   Local Socket  : 127.0.0.1..1024
                                Foreign Socket: 0.0.0.0..0
BPXOINIT 00000019   Listening   Local Socket  : 127.0.0.1..1024
                                Foreign Socket: 127.0.0.1..1025
TCPCS    00000016   Established Local Socket  : 127.0.0.1..1024
                                Foreign Socket: 127.0.0.1..1025
TCPCS    00000014   Established Local Socket  : 127.0.0.1..1024
                                Foreign Socket: 127.0.0.1..1025
4 TCP IPv4 connections

Active TCP IPv6 Connections:
Userid   Conn     State        Socket
FTPUNIX1 00000051 Listening    Local    ::0..21
                               Foreign  ::0..0
FTPMVS1  00000049 Listening    Local    ::0..1821
                               Foreign  ::0..0
2 TCP IPv6 connections

Active UDP Unicast IPv4  Connections:
Userid   Conn     Socket
PORTMAP  00000027 Local    0.0.0.0..111
```

```
                 Foreign  0.0.0.0..0
OSNMPD   00000030 Local    0.0.0.0..161
                 Foreign  127.0.0.1..162
MISCSRV  00000039 Local    198.11.98.124..7
                 Foreign  0.0.0.0..0
MISCSRV  0000003E Local    198.11.98.124..9
                 Foreign  0.0.0.0..0
4 UDP Unicast IPv4 connections

Active UDP Unicast IPv6  Connections:
Userid   Conn     Socket

0 UDP Unicast IPv6 connections

Active UDP Multicast IPv4  Connections:
Userid   Conn     Socket

0 UDP Multicast IPv4 connections

Active UDP Multicast IPv6  Connections:
Userid   Conn     Socket

0 UDP Multicast IPv6 connections

Active RAW Connections:
Userid   Conn     Protocol    Socket
TCPCS    00000006 IP          Local    0.0.0.0
                              Foreign  0.0.0.0
TCPCS    00000008 RAW         Local    0.0.0.0
                              Foreign  0.0.0.0
TCPCS    0000000E IP          Local    ::0
                              Foreign  ::0
3 RAW connections
Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS COUNTERS

Use this subcommand to display information about TCP/IP internal execution statistics.

Syntax

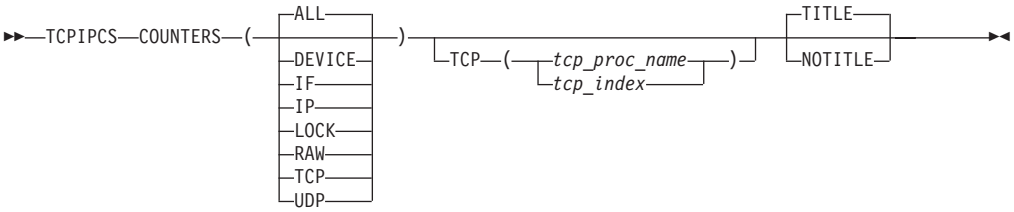

Parameters

ALL
Display all statistics. This is the default.

DEVICE
Display only device statistics.

IF Display only IF layer statistics.

IP Display only IP layer statistics.

LOCK
Display only lock statistics.

RAW
Display only RAW layer statistics.

TCP
Display only TCP layer statistics.

UDP
Display only UDP layer statistics.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS COUNTERS subcommand for IP UDP

The following is sample output of the TCPIPCS COUNTERS subcommand for IP UDP.

```
TCPIPCS COUNTERS (IP UDP)
 Dataset: SYS1.DUMP00
 Title:   LINKDOWN


 The address of the TSAB is: 15136000

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts          Status

 15136040  1 TCPCS     V1R7    1511E000 1511E0C8 002F 9FFF767F 00000000  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V1R7  found

 ================================================================================

 Analysis of Tcp/Ip for TCPCS.  Index: 1


 IP Statistics

   nonbat.............. 3
   batch............... 0
   batnum.............. 0
   nonrsm.............. 0
   batrsm.............. 0
   rsmnum.............. 0
   rteadd.............. 28
   rtedel.............. 0
   rteinc.............. 11
   rtedec.............. 8
   rtpadd.............. 14
   rtpdel.............. 0
   rtechg.............. 86
   trredr.............. 0
   trsusp.............. 0
   trsust.............. 0
   dupfrg.............. 0
   dataadj1............ 0
   dataadj2............ 0

 IP6 Statistics

   nonbat.............. 32
   batch............... 0
   batnum.............. 0
   nonrsm.............. 0
   batrsm.............. 0
```

```
  rsmnum..............  0
  rteadd..............  0
  rtedel..............  0
  rteinc..............  0
  rtedec..............  0
  rtpadd..............  11
  rtpdel..............  0
  rtechg..............  43
  trredr..............  0
  trsusp..............  0
  trsust..............  0
  dupfrg..............  0
  dataadj1............  0
  dataadj2............  0
  lifdel..............  0
  noreclaim...........  0
  hdrpullup...........  0
  dadfailtot..........  0
  dadfailll...........  0
UDP Statistics

  rd..................  7
  rdnum...............  7
  batch...............  0
  nonbat..............  7

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS DUAF

Use this subcommand to display a summary of each dispatchable unit control block (DUCB). Each entry in the dispatchable unit allocation table (DUAT) points to a DUCB. The DUAT entry contains the status of the DUCB and identifies the ASID with which the DUCB is associated. If no parameters are specified, the output contains a summary of the DUAT, followed by a summary of each DUCB.

The status of each DUCB is abbreviated as follows:

Ab The DUCB has ABENDed.

Iu The DUCB is in use.

Re The DUCB is in resume state.

Su The DUCB has been suspended.

The DUCB status might be followed by the recovery stack. There is one line for each register save area (RSA) found in the DUCB (and its DUSA extension, if present). The address of each RSA, its previous pointer, its next pointer, and the module name are shown.

A register save area displayed as `RSA*` indicates that the RSA is not in the active chain. If all RSAs are shown like this, the DUCB is not in use.

Syntax

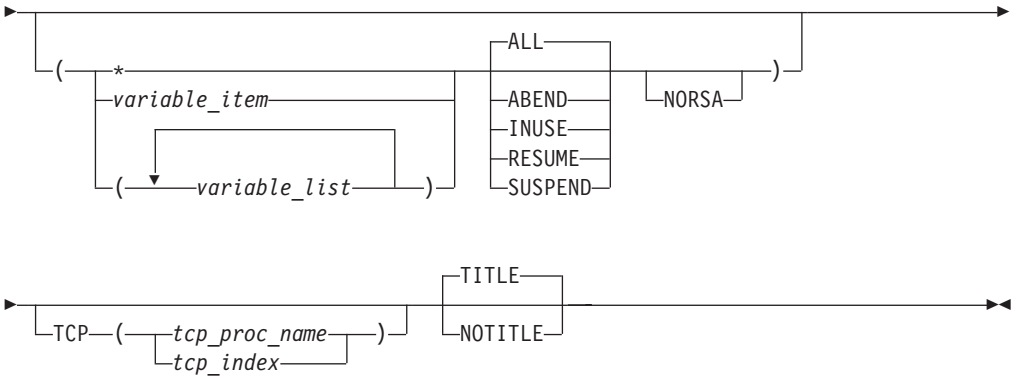

Parameters

If no parameters are specified, all active DUCBs are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

jobname
Displays only the DUCBs with this job name. The job name can be a TCP/IP application name or a stack name. Must contain from 1-8 characters.

DUCB_address
Displays the DUCB with this address. An IPCS symbol name can be specified for the address. The address is specified as 1-8 hexadecimal digits. If an address begins with characters A–F, prefix the address with a 0 to avoid the address being interpreted as a symbol name or as a character string.

DUCB_index
Displays this DUCB with this index. The index is a hexadecimal number containing one to four digits. The lowest index is 0.

ASCB_address
Displays the DUCB with this address space control block (ASCB) address. An IPCS symbol name can be specified for the address. The address is specified as 1-8 hexadecimal digits. If an address begins with characters A–F, prefix the address with a 0 to avoid the address being interpreted as a symbol name or as a character string.

ASID_number
Displays the DUCB with this ASID. The ASID is a hexadecimal number containing one to four digits.

In addition to the variable parameters described above, you can specify the following keyword parameters:

ALL
Display information for all active DUCBs. This is the default.

ABEND
Display only information for DUCBs that ABENDed.

INUSE
Display only information for DUCBs currently being used

RESUME
Display only information for DUCBs that are resumed.

SUSPEND
Display only information for DUCBs that are suspended.

NORSA
Do not display the contents of the DUCBs' register save areas (RSA). By default, the RSA contents are displayed.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {ALL, ABEND, INUSE, RESUME, SUSPEND}, only the last one is used.

Sample output of the TCPIPCS DUAF subcommand

The following example is a sample output of the TCPIPCS DUAF subcommand:

```
TCPIPCS DUAF( (0876C000  0B) INUSE )
 Dataset: IPCS.A594094.DUMPK
 Title:   TCPCS    V2R10: Job(USER15  ) EZBITRAC(HTCP50A 99.266)+
          000304 S0C4/00000004 TCB P=0029,S=000E,H=0019

 The address of the TSAB is: 08D138C0

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 08D13900  1 TCPCS     V2R10   0885A000 0885A0C8 0029 9FFFFF7F  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPCS.  Index: 1

 Dispatchable Unit Summary


 INDEX    DUAE     DUCB     DUSA     ASCB     ASID JOBNAME  ABEND    STATUS

 10000003 08859040 0876C000 0876C100 00FB7080 0029 TCPCS    00000000 Iu
 RSA  0876C3F8  Prev 00005D98  Next 0876C8C0  Mod EZBIEOER
 RSA* 0876C8C8  Prev 0876C3F8  Next 00000000  Mod EZBITSTO
  1384 bytes were used

 1000000B 08859080 08784000 08784100 00FB7980 0019 USER15   000C4000 Ab Iu
 RSA  087843F8  Prev 09BB9798  Next 087846B8  Mod EZBPFSOC
 RSA  087846C0  Prev 087843F8  Next 08784988  Mod EZBPFOPN
 RSA  08784990  Prev 087846C0  Next 08784DB0  Mod EZBUDSTR
 RSA  08784DB8  Prev 08784990  Next 087855A8  Mod EZBITRAC
  4536 bytes were used

 82 DU control blocks were found
 12 DU control blocks were in use
 0 DU control blocks were suspended
 0 DU control blocks were resumed
 1 DU control blocks had abended
 2 DU control blocks were formatted
```

```
The maximum DUCB size found is 4536 bytes

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS DUCB

Use this subcommand to display the contents of each dispatchable unit control block (DUCB). Each entry in the dispatchable unit allocation table (DUAT) points to a DUCB. The DUAT entry contains the status of the DUCB and identifies the ASID with which the DUCB is associated. The DUAT is summarized in the output. The contents of each DUCB are then displayed, followed by each DUSA for the DUCB. The first dispatchable unit stack area (DUSA) is followed by information from each register save area (RSA). Each register from the RSA is listed, showing its address and offset from the other registers in the register save area. The address of the parameter list (pointed to by R1) and the first five words at that address are also given. Each RSA is formatted. The recovery stack is also displayed.

Syntax

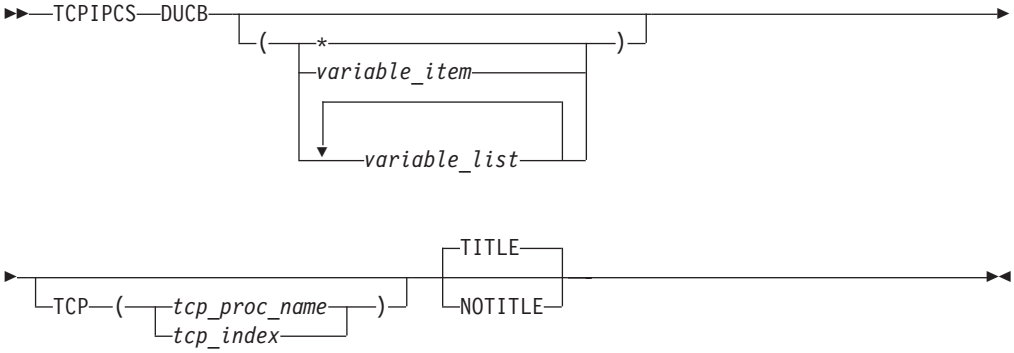

Parameters

If no parameters are specified, all DUCBs are displayed.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

jobname
Displays only the DUCBs with this job name. The job name can be a TCP/IP application name or a stack name. Must contain from 1-8 characters.

DUCB_address
Displays the DUCB with this address. An IPCS symbol name can be specified for the address. The address is specified as 1-8 hexadecimal digits. If an address begins with digit A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

DUCB_index
Displays this DUCB with this index. The index is a hexadecimal number containing one to four digits. The lowest index is zero.

ASCB_address
Displays the DUCB with this address space control block address (ASCB). An IPCS symbol name can be specified for the address. The address is specified as 1-8 hexadecimal digits. If an address begins with digit A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

asid_number
Displays the DUCB with this ASID. The ASID is a hexadecimal number containing one to four digits.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS DUCB subcommand

In the following sample, some lines have been deleted in order to shorten the sample. Deleted lines are indicated with the following:

.
.
.

The following is sample output of the TCPIPCS DUCB subcommand:

```
TCPIPCS DUCB
 Dataset: IPCS.R8A0723.RASDUMP
 Title:   EZRPE005
 The address of the TSAB is: 098221F0
 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status
 09822230  1 TCPCS     V1R5   08E85000 08E850C8 001E 9FFF7E7F  Active
 098222B0  2 TCPCS2    V1R5    08937000 089370C8 01F6 9FFF7E7F  Active
    2 defined TCP/IP(s) were found
    2 active  TCP/IP(s) were found
    2 TCP/IP(s) for CS     V1R5  found
 ================================================================================
 Analysis of Tcp/Ip for TCPCS.  Index: 1
 DUCB Detail Analysis
 Dispatchable Unit Allocation Table: 08E83010
 +0000  DUAT0      EYE..... DUAT      NEXT..... 00000000
 +0018  DUAE0      DUCB..... 08D8D010  FLAGS.... 0014001E
 +0020  DUAE1      DUCB..... 08D90000  FLAGS.... 0034001E
 +0028  DUAE2      DUCB..... 08D93000  FLAGS.... 0050001E
 +0030  DUAE3      DUCB..... 08D96000  FLAGS.... 0014001E
 +0038  DUAE4      DUCB..... 08D99000  FLAGS.... 5490001E
 .
 .
 .
 +0280  DUAE77     DUCB..... 08E74000  FLAGS.... 00000000
 +0288  DUAE78     DUCB..... 08E77000  FLAGS.... 00000000
 +0290  DUAE79     DUCB..... 08E7A000  FLAGS.... 00000000
 +0298  DUAE80     DUCB..... 08E7D000  FLAGS.... 00000000
 +02A0  DUAE81     DUCB..... 08E80000  FLAGS.... 00000000
 Dispatchable Unit Control Block: DUCB0
  EZBDUCB: 08D8D010
  +0000  DUCB_EYE.................. DUCB
  +0004  DUCB_LENGTH............... 0100
  +0006  DUCB_VERSION.............. 0002
  +0008  DUCB_TOKEN................ 08D8D010  0014001E  10000000  00000000
  +0018  DUCB_DUSA................. 08D8D110
  +001C  DUCB_AVAIL_CHAIN.......... 00000000
  +0020  DUCB_DUAEP................ 08E83028
  +0026  DUCB_ASID................. 001E
  +0028  DUCB_ASCB................. 00FA4400
  +002C  DUCB_ATCB................. 007EC920
  +0030  DUCB_ITCVT................ 08E853C8
  +0034  DUCB_LOCKSHELDCOUNT....... 00000000
  +0038  DUCB_LOCKS_TABLE.......... 08D8D194
  +003C  DUCB_LOCKS_SUSPENDED...... 00000000
  +0040  DUCB_LOCKS_SUSPENDED_NEXT. 7FFAFAF1
  +0044  DUCB_SUSPENDTOKEN......... 00000000  40000000
  +004C  DUCB_JOBNAME.............. TCPCS
  +0054  DUCB_TSAB................. 098221F0
  +0058  DUCB_TSEB................. 09822230
  +005C  DUCB_TSDB................. 08E85000
  +0060  DUCB_TSDX................. 08E850C8
  +0064  DUCB_TCP_ASCB............. 00FA4400
  +0068  DUCB_STREAMHEAD........... 00000000
  +006C  DUCB_OSI.................. 00000000
  +0070  DUCB_CREATE_FLAGS......... 00000000
  +0074  DUCB_CID.................. 00000000
```

```
 +0078  DUCB_PORT................. 0000
 +007A  DUCB_IPADDR_LEN........... 00
 +007C  DUCB_IPADDR............... 00010002  00030004  00050006  00070008
 +008C  DUCB_TRR_PTR.............. 08D8D128
 +0090  DUCB_ESTAEX_TOKEN......... 00000004
 +0094  DUCB_ERRORCODE............ 00000000
 +0098  DUCB_REASONCODE........... 00000000
 +009C  DUCB_ABEND_FLAGS.......... 000F0000
 +00A0  DUCB_DUMP_ECB............. 00000000
        DUCB_EXP_SAVE.............
 +00B8  08D8D110  00000000  00000000  88E8865C  08F671E8  7F4FC488  08D8F75C  08D8D010  08D90118  08D8F3A0  000001A8  08D8D010
 +00E8  08D8F214  08E850C8  08D8FDA8  08D8FA30  0923E21C  88E8862C

Dispatchable Unit Stack Area: DUSA0@1
 EZBDUSA: 08D8D110
 +0000  DUSA_EYE....... DUSA
 +0004  DUSA_NEXTDUSA.. 08CFF010
 +0008  DUSA_DUCB...... 08D8D010
 +000C  DUSA_START..... 08D8D450
 +0010  DUSA_LAST...... 08D90000
 +0014  DUSA_NEXTAVAIL. 08D8D540
Register Save Area: RSA0@1
Module: EZBTIWAT
 EZBRSA: 08D8D458
 +0000  RSA_DUSA. 08D8D110  RSA_PREV. 0A301890  RSA_NEXT. 00000000  RSA_R14.. FEFEFEFE  RSA_R15.. FEFEFEFE  RSA_R0... FEFEFEFE
 +0018  RSA_R1... FEFEFEFE  RSA_R2... FEFEFEFE  RSA_R3... FEFEFEFE  RSA_R4... FEFEFEFE  RSA_R5... FEFEFEFE  RSA_R6... FEFEFEFE
 +0030  RSA_R7... FEFEFEFE  RSA_R8... FEFEFEFE  RSA_R9... FEFEFEFE  RSA_R10.. FEFEFEFE  RSA_R11.. FEFEFEFE  RSA_R12.. FEFEFEFE
 +0050  RSA_AR13. FEFEFEFE  RSA_AR14. FEFEFEFE  RSA_AR15. FEFEFEFE  RSA_AR0.. FEFEFEFE  RSA_AR1.. FEFEFEFE  RSA_AR2.. FEFEFEFE
 +0068  RSA_AR3.. FEFEFEFE  RSA_AR4.. FEFEFEFE  RSA_AR5.. FEFEFEFE  RSA_AR6.. 08D773D0  RSA_AR7.. 7EFEFEFE  RSA_AR8.. 08E85084
 +0080  RSA_AR9.. 08E85238  RSA_AR10. 08E8523C  RSA_AR11. 08E85350  RSA_AR12. 08E85358
Dynamic Area of RSA0@1
Module: EZBTIWAT
 08D8D458  08D8D110  0A301890  00000000  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | .QJ............................. |
    +0020  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | ................................ |
    +0040  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | ................................ |
    +0060  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  08D773D0  7EFEFEFE  08E85084  | .....................P.}=....Y&d |
.
.
.
    +0FC0  FEFEFEFE  FEFEFEFE  08D8E620  08D8E620  FEFEFEFE  FEFEFEFE  FEFEFEFE  7F207498  | .........QW..QW.............."..q |
    +0FE0  09822230  FEFEFEFE  08E850C8  08D8D010  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | .b.......Y&H.Q}................. |
Dispatchable Unit Stack Area: DUSA0@2
 EZBDUSA: 08CFF010
 +0000  DUSA_EYE....... DUSA
 +0004  DUSA_NEXTDUSA.. 00000000
 +0008  DUSA_DUCB...... 08D8D010
 +000C  DUSA_START..... 08CFF028
 +0010  DUSA_LAST...... 08D04000
 +0014  DUSA_NEXTAVAIL. 08CFF028
Register Save Area: RSA0@2
Module: EZBCTRCD
 EZBRSA: 08CFF030
 +0000  RSA_DUSA. 08CFF010  RSA_PREV. 08D8F3A0  RSA_NEXT. 08CFF1D0  RSA_R14.. FEFEFEFE  RSA_R15.. FEFEFEFE  RSA_R0... FEFEFEFE
 +0018  RSA_R1... FEFEFEFE  RSA_R2... FEFEFEFE  RSA_R3... FEFEFEFE  RSA_R4... FEFEFEFE  RSA_R5... FEFEFEFE  RSA_R6... FEFEFEFE
 +0030  RSA_R7... FEFEFEFE  RSA_R8... FEFEFEFE  RSA_R9... FEFEFEFE  RSA_R10.. FEFEFEFE  RSA_R11.. FEFEFEFE  RSA_R12.. FEFEFEFE
 +0050  RSA_AR13. 00000000  RSA_AR14. FEFEFEFE  RSA_AR15. 08D8FDD4  RSA_AR0.. 08D8FDA8  RSA_AR1.. 08D8D010  RSA_AR2.. FEFEFEFE
 +0068  RSA_AR3.. 000263D8  RSA_AR4.. 01FF000C  RSA_AR5.. 000003C4  RSA_AR6.. 00000048  RSA_AR7.. 00000004  RSA_AR8.. 00026795
 +0080  RSA_AR9.. 08D8FA30  RSA_AR10. 08D8F75C  RSA_AR11. FEFEFEFE  RSA_AR12. FEFEFEFE
Dynamic Area of RSA0@2
Module: EZBCTRCD
 08CFF030  08CFF010  08D8F3A0  08CFF1D0  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | ..0..Q3...1}.................... |
    +0020  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | ................................ |
    +0040  FEFEFEFE  FEFEFEFE  88E88744  88E88AB8  00000000  FEFEFEFE  08D8FDD4  08D8FDA8  | ........hYg.hY...........Q.M.Q.y |
    +0060  08D8D010  FEFEFEFE  000263D8  01FF000C  000003C4  00000048  00000004  00026795  | .Q}........Q.......D..........n |
    +0080  08D8FA30  08D8F75C  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  03C40010  3001012D  | .Q...Q7*.................D...... |
    +00A0  B64C8AC7  14F03740  001E001E  001E02F0  007EC920  892336FA  E3C3D7C3  E2404040  | .<.G.0. .......0.=I.i...TCPCS    |
    +00C0  00000000  0000020A  00000000  00000000  00000000  00000000  00000000  00000000  | ................................ |
    +00E0  00000000  00000002  00000000  00000000  00000000  00000000  00000000  00000000  | ................................ |
    +0100  00000000  00000000  00000000  00000000  00000000  00000000  012892A0  00000000  | ..........................k..... |
    +0120  04FEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | ................................ |
    +0140  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  FEFEFEFE  | ................................ |
    +0160  00000000  08D8F75C  892336FA  08CFF190  08D8F3A0  08D8D010  08D8D128  08D8F214  | .....Q7*i.....1..Q3..Q}..QJ..Q2. |
    +0180  08CFF010  08D8FDA8  08CFF1D0  0923E21C  88E8862C  08CFF030  00000008  08CFF010  | ..0..Q.y..1}..S.hYf...0.......0. |

TCPIP Recovery Routine
 DTRR: 08D8D128
 +0000  TRR_CURRENT_INDEX. 00000000
 +0004  TRR_MAX_INDEX..... 00000005
         -- Array elements --
 +0008  TRR_ROUTINE....... 08E8875E
 +000C  TRR_PARM.......... 00000000
 +0010  TRR_DATA.......... 80000000
 +0014  TRR_REGS.......... 08D8D5B8
 +0018  TRR_DUMPPARM...... 00000000
 +001C  TRR_ROUTINE....... 08E8875E
 +0020  TRR_PARM.......... 00000000
 +0024  TRR_DATA.......... 80000000
 +0028  TRR_REGS.......... 08D8E358
 +002C  TRR_DUMPPARM...... 00000000
 +0030  TRR_ROUTINE....... 08E8875E
 +0034  TRR_PARM.......... 00000000
```

```
 +0038  TRR_DATA..... 80000000
 +003C  TRR_REGS..... 08CFF190
 +0040  TRR_DUMPPARM..... 00000000
 +0044  TRR_ROUTINE..... 00000000
 +0048  TRR_PARM..... 00000000
 +004C  TRR_DATA..... 00000000
 +0050  TRR_REGS..... 00000000
 +0054  TRR_DUMPPARM..... 00000000
 +0058  TRR_ROUTINE..... 00000000
 +005C  TRR_PARM..... 00000000
 +0060  TRR_DATA..... 00000000
 +0064  TRR_REGS..... 00000000
 +0068  TRR_DUMPPARM..... 00000000
        -- End of array   --

82 DUCBs were found
 2 DUCBs were formatted
Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS FRCA

Use this subcommand to display information about the Fast Response Cache Accelerator (FRCA) connections or about cached objects.

Syntax

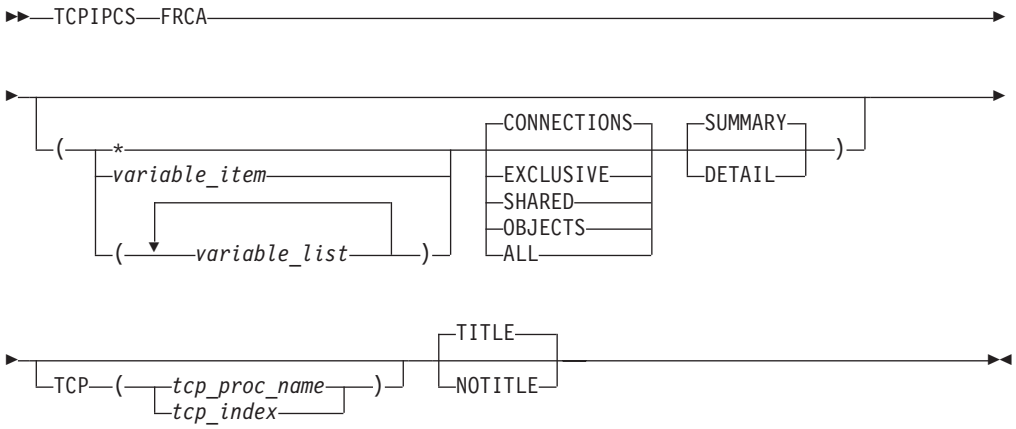

Parameters

If no parameters are specified, only FRCA connections are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

ASID Displays only the FRCA information for this address space. The address space ID is 1 - 4 hexadecimal digits in length. If the ASID value begins with digit a - f or A - F, prefix the value with a 0 so that the address is not interpreted as a symbol name or as a character string.

TCB_address
Displays the FRCA connection with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

UWSE_address
Displays the FRCA shared cache with this address. An address is 1 - 8 hexadecimal digits in length. An IPCS symbol name can be specified for an address. If an address begins with digit a-f or A-F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

UWSX_address
Displays the FRCA server connection with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a - f or A - F, prefix the address with a 0 so that the address is not interpreted as a symbol name or as a character string.

jobname
Displays only the FRCA information for this job name. The job name can be a TCP/IP application name or a stack name. The job name contains 1-8 alphanumeric characters.

connection_id
Displays the FRCA information with this connection ID. An ID is specified as 1-8 hexadecimal digits.

In addition to the variable parameters described above, you can specify the following keyword parameters:

CONNECTIONS
Display only information for FRCA connections that use either a shared or an exclusive FRCA cache. CONNECTIONS is the default.

EXCLUSIVE
Display only information for connections that use an exclusive FRCA cache.

SHARED
Display only information for connections that use a shared FRCA cache.

OBJECTS
Display only information for FRCA cached objects.

ALL
Display information for all FRCA connections and cached objects.

SUMMARY
Displays the addresses of the control blocks and other data in tables. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows the contents of the control blocks.

TCP, TITLE, NOTITLE
See “Parameters” on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {CONNECTIONS, EXCLUSIVE, SHARED, OBJECTS, ALL}, only the last one is used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS FRCA subcommand

The following is sample output of the TCPIPCS FRCA subcommand:

```
1'IPCSFRCA'                                                    1  12:03:17 01/21/08
+_____________________________________________________________________________________
0{{{{ tcpipcs frca(* all) }}}}
1'IPCSFRCA'                                                    2  12:03:17 01/21/08
+_____________________________________________________________________________________
0TCPIPCS FRCA(* ALL)
 Dataset: IPCS.MVS054.IPCSFRCA.DUMP
 Title:   'IPCSFRCA'
-The address of the TSAB is: 15927000
0Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts          Status
015927040  1 TCPCS     V1R10   1590A000 1590A0C8 002A 9FFF767F 00000000  Active
0   1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found
0   1 TCP/IP(s) for CS     V1R10 found
0================================================================================
0Analysis of Tcp/Ip for TCPCS.  Index: 1
-FRCA Server Connections
0Uwsx@    Tcb@     Cache@      References Flags Uwse@
 1642D728 7F0FD410 00000000             3 4080  16649368
 16427118 7F108390 00000000             2 4080  16649368
 1630C598 7F107F10 16436560             3 4000  00000000
0  Bucket# Uwco@       References CsmObj@         Length Key
       754 16436860             1 15785290         2,976 /index.html
0FRCA Shared Cache
0Uwse@    Cache@   References Asid
 16649368 1681F0F8          2 36
0  Bucket# Uwco@       References CsmObj@         Length Key
       754 16443240             1 15783810         2,975 /index.html
       871 164431A0             1 15785310         2,977 /index2.html
0FRCA Client Connections
0Uwcx@    Tcb@     Server@  Object@  Flags
 7F108AA8 7F108810 1642D728 00000000 10
0Analysis of Tcp/Ip for TCPCS completed
 {{{{  }}}}
 {{{{  }}}}
 {{{{  }}}}
 {{{{  ***** END OF OUTPUT ***** }}}}
```

TCPIPCS HASH

Use this subcommand to display information about the structure of TCP/IP hash tables.

Syntax

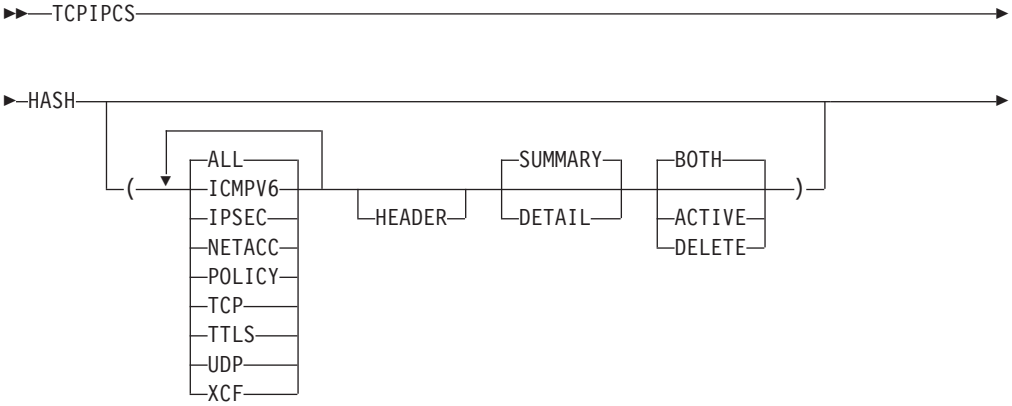

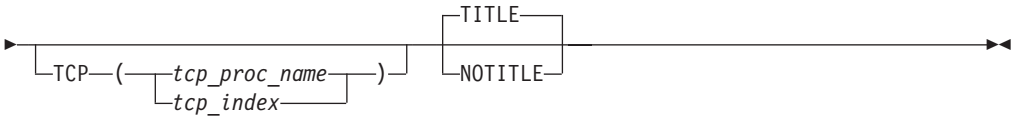

Parameters

ALL
Display structure of all TCP/IP hash tables. ALL is the default.

ICMPV6
Display only the structure of ICMPV6 hash tables.

IPSEC
Display only the structure of IPSecurity hash tables.

NETACC
Display only the structure of NetAccess hash tables.

POLICY
Display only the structure of Service Policy hash tables.

TCP
Display only the structure of TCP hash tables.

TTLS
Display only the structure of AT-TLS hash tables.

UDP
Display only the structure of UDP hash tables.

XCF
Display only the structure of XCF hash tables.

HEADER
Display hash table header information. Not displayed by default.

SUMMARY
Displays the addresses of the control blocks and other data in tables. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows the search key values.

BOTH
Display both active and logically deleted table elements. BOTH is the default.

ACTIVE
Display only the active table elements.

DELETE
Display only the logically deleted table elements.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {ALL, ICMPV6, IPSEC, NETACC, POLICY, TCP, TTLS, UDP, XCF}, all of the keywords that you specify are used.
- If you specify multiple keywords from the set {BOTH, ACTIVE, DELETE}, only the last one is used.

- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS HASH subcommand

The following is sample output of the TCPIPCS HASH subcommand.

```
TCPIPCS HASH ( DETAIL ALL )
 Dataset: D74L.KWDEV03A.DUMP
 Title:   ICMP HASHTAB
 The address of the TSAB is: 0999D6F8
 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status
 0999D738  1 TCPCS2    V1R5    08FE9000 08FE90C8 0013 00000000  Active
 0999D7B8  2 TCPCS     V1R5    08FB2000 08FB20C8 002D 00000000  Active
    2 defined TCP/IP(s) were found
    2 active  TCP/IP(s) were found
    2 TCP/IP(s) for CS     V1R5  found
 ================================================================================
 Analysis of Tcp/Ip for TCPCS2.  Index: 1
 TCPIP Hash Table Analysis
 Policy ID Port Table
 Hash Table Header at 7F65E008
   Instance            : 1
   Active entries      : 0
   Hash buckets        : 1,999
   User free routine   : 00000000
   Element queue       : 08FE9E48
 0 elements in Policy Id Port Table
 Table Summary:
   Active buckets      : 0
   Inactive buckets    : 0
   Unused buckets      : 1,999
   Max active q length : 0
   Max active q index  : 0
   Max active q seqnum : 0
   Max delete q length : 0
   Max delete q index  : 0
   Total seqnum        : 0


 ICMPV6         Table
 Hash Table Header at 7F699C08
   Instance            : 4
   Active entries      : 7
   Hash buckets        : 1,024
   User free routine   : 00000000
   Element queue       : 08FE9E50
 Bucket# Bucket@  Element@ Status  User@    KeyValue
       2 7F699C28 7F2F8E80 Active  0AA6BFD8 FEC00000 00000000 00000000 00000000
                                            Clock Ticks...... 00000003
                                            Tokens........... 00
                                            Token Tenths..... 00
       5 7F699C58 7F2F9100 Active  0AA6BF98 00000000 00000000 00000000 00000000
                                            Clock Ticks...... 00000008
                                            Tokens........... 00
                                            Token Tenths..... 00
       6 7F699C68 7F2F9080 Active  0AA6BFA8 00000000 00000000 00000000 00000000
                                            Clock Ticks...... 00000011
                                            Tokens........... 00
                                            Token Tenths..... 00

 7 elements in ICMPV6         Table
 Table Summary:
   Active buckets      : 6
   Inactive buckets    : 0
   Unused buckets      : 1,018
   Max active q length : 2
   Max active q index  : 6
```

```
  Max active q seqnum : 2
  Max delete q length : 0
  Max delete q index  : 0
  Total seqnum        : 7
TCP V4 Index Table
Hash Table Header at 7F528B88
  Instance            : 2
  Active entries      : 6
  Hash buckets        : 62,533
  User free routine   : 88D9523E
  Element queue       : 08FE9E48
Bucket# Bucket@  Element@ Status  User@    KeyValue
      0 7F528B88 7F507FE0 Active  7F510108 00000000 00000000 00000000
    530 7F52ACA8 7F5080C0 Active  7F51B988 00000000 00000000 00150000
  30479 7F59FC78 7F508020 Active  7F510A08 7F000001 00000000 04000000
  35181 7F5B2258 7F5080A0 Active  7F51A788 00000000 00000000 27170000
  37771 7F5BC438 7F508040 Active  7F511308 7F000001 7F000001 04000401
  40773 7F5C7FD8 7F508060 Active  7F510E88 7F000001 7F000001 04010400
6 elements in TCB V4 Index Table
Table Summary:
  Active buckets      : 6
  Inactive buckets    : 1
  Unused buckets      : 62,526
  Max active q length : 1
  Max active q index  : 0
  Max active q seqnum : 1
  Max delete q length : 0
  Max delete q index  : 0
  Total seqnum        : 8

TCP V6 Index Table
Hash Table Header at 7F2FDB88
  Instance            : 5
  Active entries      : 2
  Hash buckets        : 62,533
  User free routine   : 88D9523E
  Element queue       : 08FE9E50
Bucket# Bucket@  Element@ Status  User@    KeyValue
      0 7F2FDB88 7F2F8C80 Active  7F510588 00000000 00000000 00000000 00000000
                                           00000000 00000000 00000000 00000000
                                           00000000
    530 7F2FFCA8 7F2F8F00 Active  7F51B988 00000000 00000000 00000000 00000000
                                           00000000 00000000 00000000 00000000
                                           00000000
2 elements in TCB V6 Index Table
Table Summary:
  Active buckets      : 2
  Inactive buckets    : 0
  Unused buckets      : 62,531
  Max active q length : 1
  Max active q index  : 0
  Max active q seqnum : 1
  Max delete q length : 0
  Max delete q index  : 0
  Total seqnum        : 2

UDP DMUX V4 Table
Hash Table Header at 7F403B88
  Instance            : 3
  Active entries      : 2
  Hash buckets        : 62,533
  User free routine   : 88DB0E3C
  Element queue       : 08FE9E48
Bucket# Bucket@  Element@ Status  User@    KeyValue
      0 7F403B88 7F508000 Active  7F4F8108 00000000 0000
    529 7F405C98 7F508080 Active  7F500608 00000000 0211
```

```
2 elements in UDP DMUX V4 Table
Table Summary:
  Active buckets      : 2
  Inactive buckets    : 0
  Unused buckets      : 62,531
  Max active q length : 1
  Max active q index  : 0
  Max active q seqnum : 1
  Max delete q length : 0
  Max delete q index  : 0
  Total seqnum        : 2
UDP DMUX V6 Table
Hash Table Header at 7F203B88
  Instance            : 6
  Active entries      : 1
  Hash buckets        : 62,533
  User free routine   : 88DB0E3C
  Element queue       : 08FE9E50
Bucket# Bucket@  Element@ Status  User@    KeyValue
      0 7F203B88 7F2F8D00 Active  7F4F8808 00000000 00000000 00000000 00000000
                                           0000
1 elements in UDP DMUX V6 Table
Table Summary:
  Active buckets      : 1
  Inactive buckets    : 0
  Unused buckets      : 62,532
  Max active q length : 1
  Max active q index  : 0
  Max active q seqnum : 1
  Max delete q length : 0
  Max delete q index  : 0
  Total seqnum        : 1

UDP MULTICAST V6 Table
Hash Table Header at 7F10EB88
  Instance            : 7
  Active entries      : 0
  Hash buckets        : 62,533
  User free routine   : 88DB0E3C
  Element queue       : 08FE9E50
0 elements in UDP MULTICAST V6 Table
Table Summary:
  Active buckets      : 0
  Inactive buckets    : 0
  Unused buckets      : 62,533
  Max active q length : 0
  Max active q index  : 0
  Max active q seqnum : 0
  Max delete q length : 0
  Max delete q index  : 0
  Total seqnum        : 0
Analysis of Tcp/Ip for TCPCS2 completed
```

TCPIPCS HEADER

Use the TCPIPCS HEADER command to display information from the system dump header and, in some cases, if a DUCB has ABENDed, the DUCB is displayed. The IPCS command `STATUS System Cpu Registers Worksheet Faildata` is used to display the system dump header.

Depending on the error recovery routine, the DUCB address might or might not be available. If the DUCB address is available, the DUCB is displayed. To find DUCBs that ABENDed, use the TCPIPCS DUAF (* ABEND) command.

Syntax

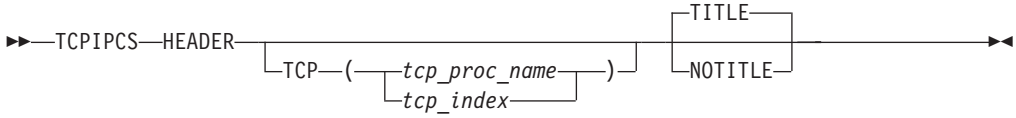

Parameters

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS HEADER subcommand

The following is sample output of the TCPIPCS HEADER subcommand:

```
TCPIPCS HEADER
 Dataset: IPCS.MV21381.DUMPA
 Title:   SLIP DUMP ID=TC


 The address of the TSAB is: 13391BC0

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 13391C00  1 TCPSVT    V2R10   1323B000 1323B0C8 07DE 04041405  Active
 13391C80  2 TCPSVT2   V2R10   00000000 00000000 07E8 00000000  Down Stopping
 13391D00  3 TCPSVT1   V2R10   12FC3000 12FC30C8 0080 94FF755F  Active
 13391D80  4 TCPSVT3   V2R10   00000000 00000000 0059 00000000  Down Stopping

    4 defined TCP/IP(s) were found
    2 active  TCP/IP(s) were found

    4 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPSVT.  Index: 1

                                  STATUS SUBCOMMAND


                              MVS Diagnostic Worksheet

 Dump Title: SLIP DUMP ID=TC

 CPU Model 9672 Version AC Serial no. 041018 Address 00
 Date: 03/22/2000     Time: 07:36:57.297123 Local

 Original dump dataset: SYS1.DUMP93

 Information at time of entry to SVCDUMP:

 HASID 000B  PASID 000B  SASID 000B  PSW 440C0000 81584B1C

 CML ASCB address 00000000  Trace Table Control Header address 7F45D000

 Dump ID: 007
 Error ID: N/A

 SDWA address N/A

....

CPU STATUS:
```

```
 PSW=440C0000  81584B1C  (RUNNING IN PRIMARY, KEY 0,  AMODE 31, DAT ON)
     DISABLED FOR I/O EXT
   ASID(X'000B') 01584B1C. IEANUC09.IEAVEDS0+1C IN READ ONLY NUCLEUS
  ASCB11 at FBD700, JOB(WLM), for the home ASID
  ASXB11 at 7FDFA0 and TCB11M at 7FB440 for the home ASID
  HOME ASID: 000B PRIMARY ASID: 000B SECONDARY ASID: 000B

  GPR VALUES
      0-3  00000001  0288E01C  00000C38  00000008
      4-7  007FB440  007FFC10  007F6A68  00FBD700
      8-11 00000000  01584B00  015AD820  007FEE48
     12-15 00000000  00000000  80FDE336  81584B18
...
```

TCPIPCS HELP

Use this subcommand to display TCPIPCS usage and syntax information.

Syntax

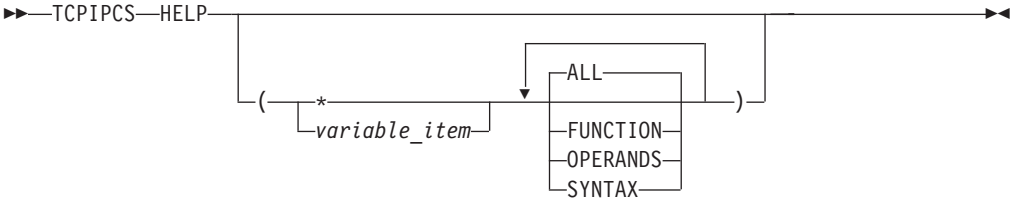

Parameters

If no parameters are specified, the function, operand, and syntax information are displayed for all TCPIPCS commands.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the TCPIPCS subcommand names.

In addition to the variable parameters described above, you can specify the following keyword parameters:

ALL
Display information for all TCPIPCS commands. ALL is the default.

FUNCTION
Display only function information.

OPERANDS
Display only operand information.

SYNTAX
Display only syntax information.

Rule: If you specify multiple keywords from the set {ALL, FUNCTION, OPERANDS, SYNTAX}, all of the keywords that you specify are used.

Sample output of the TCPIPCS HELP subcommand

The following is sample output of the TCPIPCS HELP subcommand:

```
tcpipcs help (config function)

 Function:

    The TCPIPCS command displays selected information about a specific
```

```
   TCP/IP address space.

  CONFIG - Produce device configuration report.

   Function:

     Display information about device, physical, and logical interfaces

   Syntax:

     TCPIPCS CONFIG(<{SUMMARY|DETAIL}>)

   Operands:

     SUMMARY - Display summary report.

     DETAIL - Display summary and interface cross-reference reports.

***
```

TCPIPCS IPSEC

Use this subcommand to display information about IP security filters or tunnels, IP security translated ports, or defensive filters.

Syntax

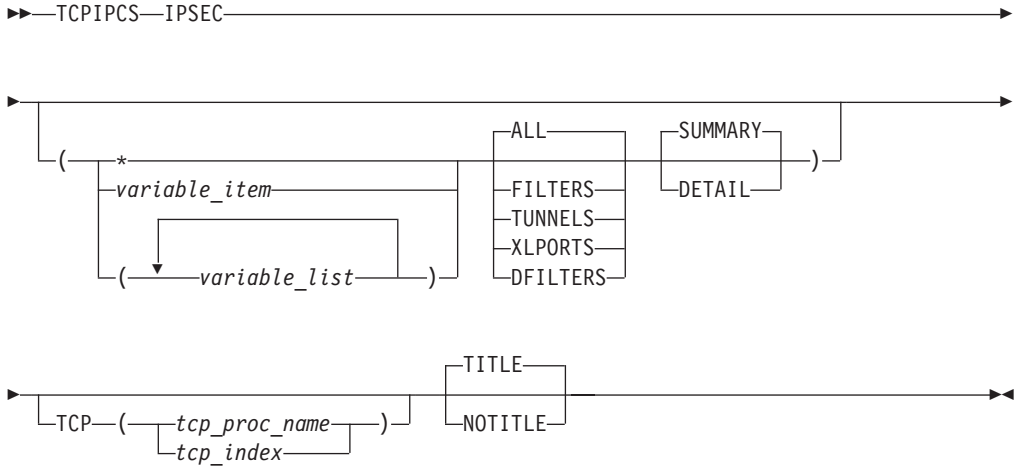

Parameters

If you do not specify any parameters, all IP security filters, tunnels, NAT translated ports, and all defensive filters are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
The following variable parameters can be repeated up to 32 times, separated by a blank space, within parentheses:

filter_address
Displays the IP security filter or defensive filter that has this address. An address is specified as 1 - 8 hexadecimal digits. An IPCS symbol name can

be specified for an address. If an address begins with a - f or A - F, prefix the address with a 0 so that the address is not interpreted as a symbol name or as a character string.

tunnel_address
Displays the IP security tunnel that has this address. An address is specified as 1 - 8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with a - f or A - F, prefix the address with a 0 so that the address is not interpreted as a symbol name or as a character string.

source_IP_address
Displays the IPSecurity NAT SourceIP table entry with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a-f or A-F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

translated_port_address
Displays the IPSecurity NAT Port Translation table entry with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a-f or A-F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

In addition to the variable parameters previously described, you can specify the following keyword parameters:

ALL
Display information for IP security filters, tunnels, NAT traversal remote port translations, and defensive filters. ALL is the default.

FILTERS
Display only information for IP security filters.

TUNNELS
Display only information for IP security tunnels.

XLPORTS
Display only information for IP security NAT-translated ports.

DFILTERS
Display only information for defensive filters.

SUMMARY
Displays the addresses of the control blocks and other data in tables. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows the contents of the control blocks.

TCP, TITLE, NOTITLE
See “Parameters” on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {ALL, FILTERS, TUNNELS, XLPORTS, DFILTERS}, only the last one is used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Restriction: The TCPIPCS IPSEC subcommand works only on stacks configured for IP security.

Sample output of the TCPIPCS IPSEC subcommand

The following is sample output of the TCPIPCS IPSEC subcommand:

```
TCPIP Ipsecurity Analysis

IPSEC on ZIIP No

FWE_GDA   at 7F283430
FILTER_DA at 7F172C10
DEFENSIVE_FILTER_DA at 7F172BB0
TUNNEL_DA at 7F172530
ENCRYP_DA at 7F282890

Filter set active : Policy
Filter logging : No
Pre-decap filtering : No

Defense Filter Mode : Active
  IPv4 Defensive filter inbound list: 00000000
  IPv4 Defensive filter inbound count: 0
  IPv4 Defensive filter outbound list: 00000000
  IPv4 Defensive filter outbound count: 0
  IPv6 Defensive filter inbound list: 00000000
  IPv6 Defensive filter inbound count: 0
  IPv6 Defensive filter outbound list: 00000000
  IPv6 Defensive filter outbound count: 0

IPv4 Filters

Filter@  Action  SPrt1 SPrt2 DPrt1 DPrt2 Protocol
                 Src@
                 Dst@
7C553110 Permit    500     0   500     0  17 (UDP)
                 0.0.0.0/0
                 0.0.0.0/0
7B8DBD90 Permit      0     0   623     0   6 (TCP)
                 197.11.107.1
                 197.11.236.12


IPv6 Filters

Filter@  Action  SPrt1 SPrt2 DPrt1 DPrt2 Protocol
                 Src@
                 Dst@
7B8D1610 Permit    500     0   500     0  17 (UDP)
                 ::0/0
                 ::0/0
7DA22D90 Permit    623     0     0     0   6 (TCP)
                 2000:197:11:235::101:0:1
                 2000:197:11:107::1


IPv4 Tunnels

Tunnel@  Policy   Format   Name
         Src@
         Dst@
7C514010 000000A5 00000033 Y 1589 DVA-linux
         197.11.235.9
         16.11.16.126


IPv6 Tunnels
```

```
Tunnel@  Policy   Format   Name
         Src@
         Dst@
7E5AF010 0000014A 0000000C M 1 IPMVAospfAH03
         ::0
         ::0
```

TCPIPCS LOCK

Use this subcommand to scan the dump for information about the current locks that are defined and held.

Only nonzero statistics are reported.

Tip: The DUCB lock table entries might conflict with the lockword counters. This is because DUCB lock table entries and lockword counters are not updated in one operation, therefore they can be out of sync. At the time the dump was obtained, the lockword counters might have been updated, but the DUCB has not yet been updated.

Syntax

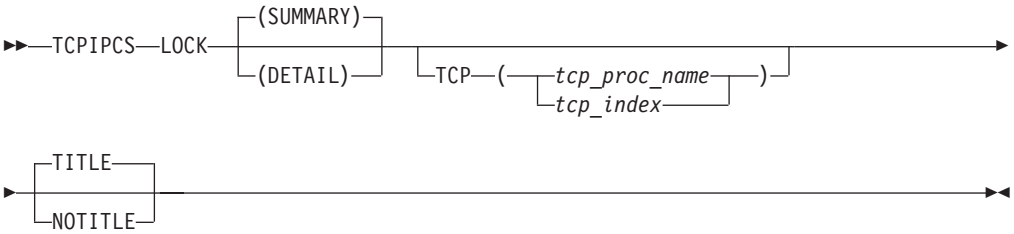

Parameters

SUMMARY
Displays each level of each class of lock, the total number of DUCBs found, and a cross-reference for each lock being used. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows lock information for each DUCB.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS LOCK subcommand

The following is sample output of the TCPIPCS LOCK subcommand:

```
TCPIPCS LOCK (DETAIL)
 Dataset: IPCS.A594094.DUMPM
 Title:   TCPSVT   V3R10: Job(TCPSVT  ) EZBITSTO(HTCP50A 99.281)+
          00077A S4C5/74BE2500 SRB P=0051,S=0051,H=0051
 ...
 ItCvt: 12B573C8, Class_Count: 12, Level_Count: 34, Table_Size: 616

 Lock statistics at 12E7B208

   Class 2 at 12E7B2E8 for 2 levels
     Level 0201 ITSTOR_QUE
```

```
      Suspension - Srb :           1,601
      Delays     -     :             239
  ...
  Class 6 at 12E7B478 for 4 levels
    Level 0602 TCB
      Suspension - Srb :             146
      Suspension - Tcb :              33
  ...

Ix   Ducb@    Lktb@    Susp@    Next@    DucbIx   Status
0002 12A62000 12A62184 00000000 00000000 10000001 Iu
  Lock Class 02: 00000001 00000002 12A62278 00000000
    Lock Level 01: 12B57CB8 C0010201 00010000 Held Excl       ITSTOR_QUE

Ix   Ducb@    Lktb@    Susp@    Next@    DucbIx   Status
072E 12B19000 12B19184 00000000 7FFAFAF1 1000003E Iu
  Lock Class 06: 00000002 00000004 12B192F0 00000000
    Lock Level 02: 7F272D38 80010602 00020100 Held Shr        TCB

50 DUCBs found
2 DUCBs held locks
0 DUCBs were waiting for locks

Lockword Cross Reference

Lock@    Ducb@    Status          Name
12B57CB8          Not Held        ITSTOR_QUE
7F272D38 12B19000 Held Shr        TCB

2 locks were referenced

Lock Class/Level Multiple Usage:

Class Level Names
   03    02 REASM
            PTREE
            MCGRP
   ...
   0C    06 SKITSSL
            TCFG_CLEANUP

Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS MAP

Use this subcommand to display a mapping of TCP/IP storage. This subcommand is useful for finding overlays and abandoned storage.

Each control block referenced is listed in order by its address. Each control block eye-catcher is shown; if none is found, a mnemonic name is given in quotation marks. The size is the number of bytes (in decimal) in the control block. The key is the storage key. The base and offset are the address of a TCP/IP control block and the offset within it that contains the CbAddr in the far left column. Multiple references can exist, so additional references are continued on a separate line.

Note: Large dumps with many control blocks can take considerable time to process.

Syntax

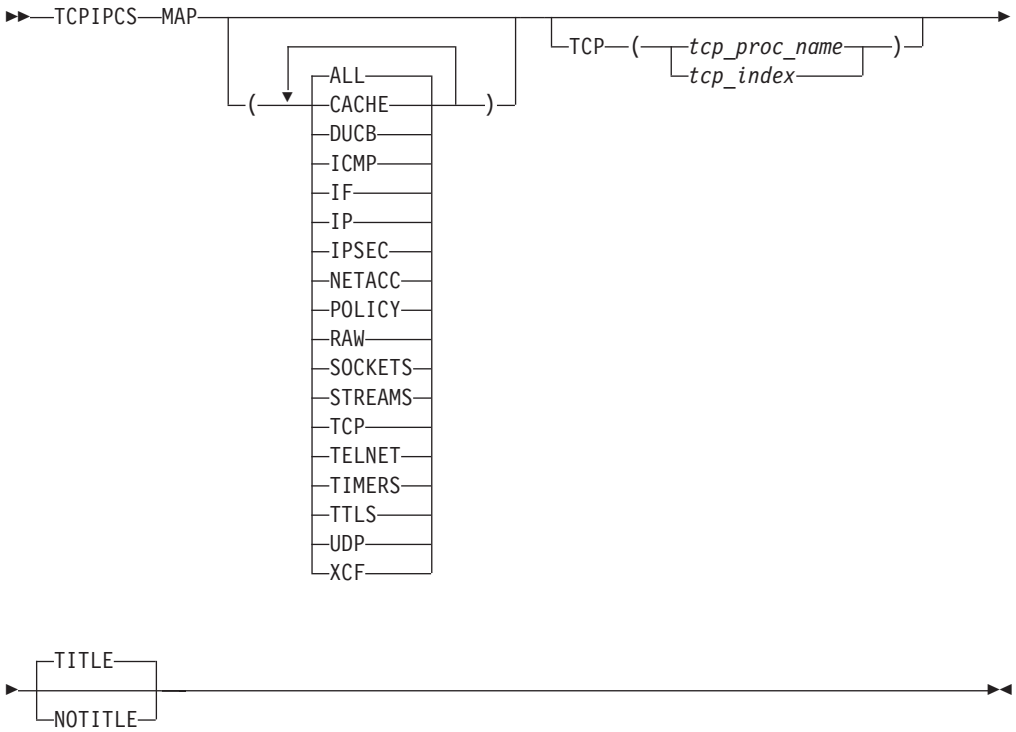

Parameters

ALL
Display storage usage information for all components.

CACHE
Display only CACHE storage usage information.

DUCB
Display only DUCB storage usage information.

ICMP
Display only ICMP storage usage information.

IF Display only IF/IP storage usage information.

IP Display only IF/IP storage usage information.

IPSEC
Display only IPSEC storage usage information.

NETACC
Display only NETACC storage usage information.

POLICY
Display only POLICY storage usage information.

RAW
Display only RAW storage usage information.

SOCKETS
Display only SOCKETS storage usage information.

STREAMS
Display only STREAMS storage usage information.

TCP
Display only TCP storage usage information.

TELNET
Display only TELNET storage usage information.

TIMERS
Display only TIMERS storage usage information.

TTLS
Only display AT-TLS storage usage information.

UDP
Display only UDP storage usage information.

XCF
Display only XCF storage usage information.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {ALL, CACHE, DUCB, ICMP, IF, IP, IPSEC, NETACC, POLICY, RAW, SOCKETS, STREAMS, TCP, TELNET, TIMERS, TTLS, UDP, XCF}, all of the keywords that you specify are used.

Sample output of the TCPIPCS MAP subcommand

The following is sample output of the TCPIPCS MAP subcommand:

```
TCPIPCS MAP
 Dataset: IPCS.MV20767.DUMPA
 Title:   VERIFY MV20758


 The address of the TSAB is: 08DD36F8

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 08DD3738  1 TCPCS     V2R10   0876E000 0876E0C8 01F7 92208100  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPCS.  Index: 1

 CbIds enclosed in quotes e.g. "CBID" are not true eyecatchers.

 Found 847 References and 1037 Cross-references

 CbAddr   CbId              Size Key    Base    +Offset
 00FCC6A0  CVT            1,280 6
 01663450 ECVT              576 6      00FCC6A0+008C
 0876B000 "ALCCSA"           96 6      08DD9328+0004
                                       08DD9368+0000
 0876B388 "CACSMM"          120 6      0876B408+0004
 0876B408 "CACSMM"          120 6      0876E5C8+0560
 0876B488 "CACSMM"          120 6      0876B688+0004
 0876B500 "CACSA "          120 6      0876B600+000C
 0876B580 "CACSA "          120 6      0876B500+000C
 0876B600 "CACSA "          120 6      0876E5C8+0218
 0876B688 "CACSMM"          120 6      0876E5C8+0568
 0876B700 "CACSA "          120 6      0876D700+000C
 0876B780 "CACSA "          120 6      0876B700+000C
 ...
 7F6E8B78 SKQU               64 6      7F6E8748+00E8
 7F6E8BB8 SKQU               64 6      7F6E8AA8+0004
```

```
7F6E8BF8 SKQP               16 6      7F6E8B78+0018
                                      7F6E8BB8+0018
7F6E8C08 SKBD               32 6      7F6E8B78+002C
7F6E8C28 SKBD               32 6      7F6E8C08+0004
7F6E8C48 SKBD               32 6      7F6E8BB8+002C
7F6E8C68 SKBD               32 6      7F6E8C48+0004
7F6E8CC8 SKSC              176 6      7F6E8008+0004
                                      7F6E8748+0060
7F6E8D88 SKRT              128 6      0876E0C8+0130

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS MTABLE

Use this subcommand to access the module tables and display the following information:

- Module entry point address
- Name
- Compile date and time
- PTF number
- Load module name

The entries are listed first in entry-point-address order, and then listed again in module-name order.

Syntax

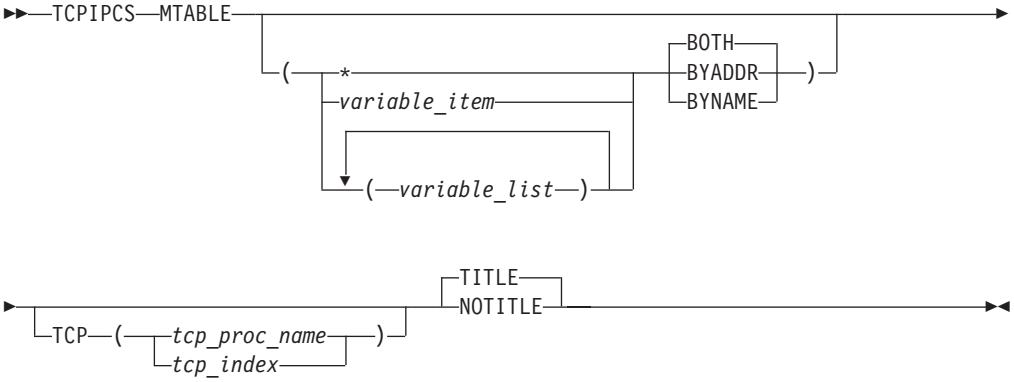

Parameters

If no parameters are specified, all displayable modules are displayed.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 –32 of the following variable parameters, each separated by a blank space, within parentheses:

address
Locates the TCP/IP module where this address appears and displays the name and offset. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address

begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

name Locates the TCP/IP module with this name. A name is specified as 1-8 characters.

In addition to the variable parameters previously described, you can specify the following keyword parameters:

BOTH Display modules sorted by address and by name.

BYADDR
Display only modules sorted by address.

BYNAME
Display only modules sorted by name.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS MTABLE subcommand

The following is a sample output of the TCPIPCS MTABLE subcommand:

```
TCPIPCS MTABLE (12DE3800 12D9B858)
 Dataset: IPCS.A594094.DUMPM
 Title:   TCPSVT   V2R10: Job(TCPSVT  ) EZBITSTO(HTCP50A 99.281)+
          00077A S4C5/74BE2500 SRB P=0051,S=0051,H=0051

 The address of the TSAB is: 12E89BB8

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 12E89BF8  1 TCPSVT    V2R10   12B57000 12B570C8 0051 9FFFFF7F  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPSVT.  Index: 1

 TCPIP Module Table Analysis

 TCMT 12B590E8 EZBITCOM Size: 00D8 Cnt: 47
 MTBL 12C23F28 EZBTIINI Size: 0CD4 Cnt: 272
 MTBL 948ACA50 EZBTZMST Size: 0134 Cnt: 24
 MTBL 94FE8470 EZBTTMST Size: 0704 Cnt: 148
 MTBL 94AA0B00 EZBTMCTL Size: 0380 Cnt: 73

 Module   Epa      Date       Time     PTF      Lmod     Asid

 EZBIFARP 12DE35D8 1999/10/15 07:01:58 HTCP50A  EZBTIINI 0051
 EZBXFINI 12D9B808 1999/10/08 00:37:29 HTCP50A  EZBTIINI 0051

 Address 12DE3800 is EZBIFARP+0228
 Address 12D9B858 is EZBXFINI+0050

 Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS POLICY

Use this subcommand to display policy information.

Syntax

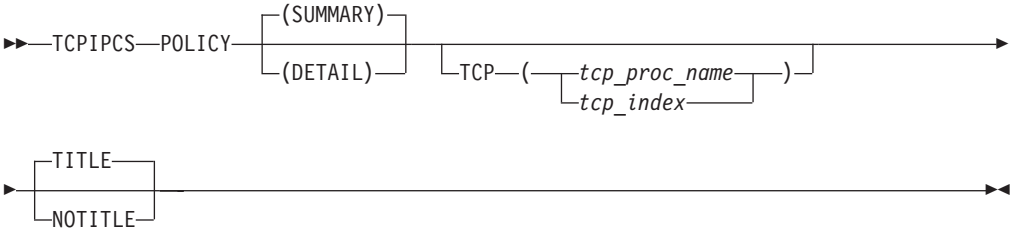

Parameters

SUMMARY
Displays the policy table addresses. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows control block contents.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS POLICY subcommand

The following is sample output of the TCPIPCS POLICY subcommand:

```
TCPIPCS POLICY TCP(1)
Dataset: IPCS.MV21046.DUMPA
Title:   BOTSWANA HUNG RUNNING PAGENT DIFFSERV SETTINGS.


The address of the TSAB is: 12EFD818

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

12EFD858  1 TCPSVT    V2R10   12EAB000 12EAB0C8 0058 9CFF755F  Active
12EFD8D8  2 TCPSVT1   V2R10   12A0F000 12A0F0C8 0069 9CFF755F  Active
12EFD958  3 TCPSVT2   V2R10   127C9000 127C90C8 07DE 9CFF755F  Active
12EFD9D8  4 TCPSVT3   V2R10   126FB000 126FB0C8 0054 9CFF755F  Active
12EFDA58  5 TCPSVT4   V2R10   12646000 126460C8 004C 9CFF755F  Active
12EFDAD8  6 TCPSVT5   V2R10   1260E000 1260E0C8 07DD 9CFF755F  Active
12EFDB58  7 TCPSVT6   V2R10   12383000 123830C8 007A 9CFF755F  Active
12EFDBD8  8 TCPSVT7   V2R10   11ECE000 11ECE0C8 07DC 9CFF755F  Active

   8 defined TCP/IP(s) were found
   8 active  TCP/IP(s) were found

   8 TCP/IP(s) for CS     V2R10 found

================================================================================

Analysis of Tcp/Ip for TCPSVT.  Index: 1


Policy Control Table at 12F54210

Intrusion Detection Main Table at 13AA6088

Service Classes:

Scentry@ Scope Tos Pri Permission Name
```

```
129455F0 Both  60  00  Allowed    paPRD-GenImp5
129454F0 Both  00  00  Allowed    padefault
129453F0 Both  E0  00  Allowed    paOSPF-1
129452F0 Both  E0  00  Allowed    paTST-1-GenImp1
129451F0 Both  C0  00  Allowed    paTST-1-GenImp2
12942B10 Both  A0  00  Allowed    paTST-1-GenImp3
12942A10 Both  80  00  Allowed    paTST-1-GenImp4
12942910 Both  60  00  Allowed    paTST-1-GenImp5
12942810 Both  40  00  Allowed    paTST-1-GenImp6
12942710 Both  20  00  Allowed    paTST-1-GenImp7
...

Policy Rules:

Prentry@ Permission Cond Level@   Cond@    Name
126C04F0 Allowed    DNF  00000000 00000000 prPRD-CBS-30001
128F2A90 Allowed    CNF  126EB590 00000000 prTST-WEB-4-80-BO
                         126C0B10 00000000
                         126C0790 126C0950
...
```

TCPIPCS PROFILE

Use this subcommand to show the active configuration information at the time of the dump, in the form of profile data set statements. This profile does not necessarily match the profile used to start TCP/IP because the startup profile might not include the dynamic changes, additions, or deletions made by using commands. All the defaults that are in effect are displayed in addition to explicit settings.

Syntax

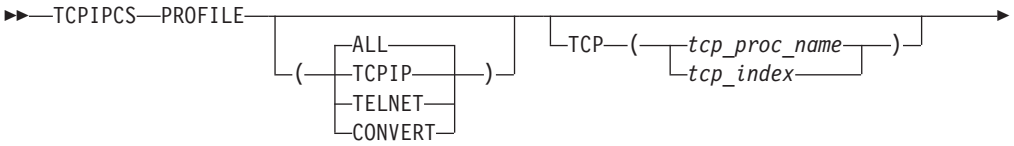

Parameters

ALL
Display all profile statements.

TCPIP
Display only TCP/IP profile statements.

TELNET
Display only Telnet profile statements.

CONVERT
Display converted profile statements. This parameter can be useful as an aid to convert older profile statements to equivalent strategic profile statements.

For TCPIP, the following are the converted profile statements.

- The DEVICE, LINK, and HOME statements for the following interface types are converted to INTERFACE statements:
  - IPAQENET

- IPAQIDIO
- VIRTUAL

In addition, corresponding changes are made to the BSDROUTINGPARMS, START, and STOP statements if present.

- The SMFPARMS statement is converted to an SMFCONFIG statement.

Rule: The SMFPARMS statement is converted to an SMFCONFIG statement even if the CONVERT parameter is not specified.

For TELNET, there are no converted profile statements.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS PROFILE subcommand

The following is sample output of the TCPIPCS PROFILE subcommand:

```
TCPIPCS PROFILE
Dataset: IPCS.TCPIPCS.DUMP
Title:   TCPIPCS PROFILE SAMPLE IPV6 CINET
The address of the TSAB is: 162E9000
Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts          Status
162E9040  1 TCPCS5    V2R1   15951000 159510C8 0032 9FFF777F 00000000  Active
162E90C0  2 TCPCS8    V2R1   14FEB000 14FEB0C8 002E 9FFF777F 00000000  Active
   2 defined TCP/IP(s) were found
   2 active  TCP/IP(s) were found
   2 TCP/IP(s) for CS     V2R1 found
================================================================================
Analysis of Tcp/Ip for TCPCS5.  Index: 1
;
; Profile generated on 2015/01/30 at 17:34:03
;
; Dump Dataset  : IPCS.TCPIPCS.DUMP
; Dump Time     : 2015/01/30 12:33:18.247759
; TCP/IP Jobname: TCPCS5
;


;
; For informational purposes, only SMFCONFIG
; will be generated in this reconstructed
; profile.
;
; Either an SMFCONFIG or an SMFPARMS statement
; is specified in an initial profile data set,
; or in a data set referenced by the
; VARY TCPIP,,OBEYFILE command.
; SMFCONFIG is the recommended way to define
; SMF processing options.
;
ARPAGE 20
AUTOLog 5
  TRMD     PARMSTRING "D=TCPDATA5"
ENDAUTOLog
DEVice VIPA4815 VIRTual 0000
LINK VIPA4815L VIRTual 0 VIPA4815
DEVice IUTSAMEH MPCPTP NOAUTORESTART
LINK TOVTAM MPCPTP IUTSAMEH IFSPEED 4500000 CHECKSUM SECCLASS 255
     NOMONSYSPLEX
DEVice MPC4115 MPCPTP NOAUTORESTART
LINK MPC4115L MPCPTP MPC4115 IFSPEED 4500000 CHECKSUM SECCLASS 255
     NOMONSYSPLEX
DEVice MPC4145 MPCPTP NOAUTORESTART
LINK MPC4145L MPCPTP MPC4145 IFSPEED 4500000 CHECKSUM SECCLASS 255
     NOMONSYSPLEX
DEVice MPC4185 MPCPTP NOAUTORESTART
```

```
LINK MPC4185L MPCPTP MPC4185 IFSPEED 4500000 CHECKSUM SECCLASS 255
     NOMONSYSPLEX
DEVice QDIO4105 MPCIPA PRIROUTER NOAUTORESTART
GLOBALCONFig NOTCPIPStatistics ECSALIMIT 0K POOLLIMIT 0K
             NOMLSCHKTERMinate NOEXPLICITBINDPORTRANGE SYSPLEXMONitor
             TIMERSECS 60 NORECOVERY NODELAYJOIN NOAUTOREJOIN
             NOMONINTERFACE NODYNROUTE SYSPLEXWLMPoll 60 ZIIP
             NOIPSECURITY NOIQDIOMULTIWRITE MAXRECS 300 IQDMULTIWRITE
             NOAUTOIQDX WLMPRIORITYQ IOPRI1 0 IOPRI2 1 IOPRI3 2 3
             IOPRI4 4 5 6 FWD
NETMONitor NONTATRCService NOPKTTRCService NOTCPCONNService SMFService
NOIPSECURITY PROFILE DVIPA
HOME
  10.81.5.5 VIPA4815L
  10.51.0.5 TOVTAM
  10.11.1.5 MPC4115L
  10.11.4.5 MPC4145L
  10.11.8.5 MPC4185L
INTERFace I4VIPA1 DEFINE VIRTUAL IPADDR 10.81.5.6
INTERFace QDIO4105I DEFINE IPAQENET CHPIDTYPE OSD PORTNAME QDIO4105 IPADDR
          172.16.1.5/24 PRIROUTER MTU 8992 READSTORAGE GLOBAL INBPERF
          BALANCED NOMONSYSPLEX NODYNVLANREG NOISOLATE NOOLM
INTERFace I4HIPERSOCKA DEFINE IPAQIDIO CHPID 9A IPADDR 172.18.1.1/32
          READSTORAGE GLOBAL SOURCEVIPAINTerface I4VIPA1 MONSYSPLEX
INTERFace VIPA6815 DEFINE VIRTUAL6 IPADDR 2001:DB8:10::81:5:5
INTERFace IUTSAMEH6 DEFINE MPCPTP6 TRLE IUTSAMEH INTFID ::51:0:5
          SECCLASS 255 NOMONSYSPLEX IPADDR 2001:DB8:10::51:0:5
INTERFace MPC6115 DEFINE MPCPTP6 TRLE MPC4115 INTFID ::11:1:5
          SOURCEVIPAINTerface VIPA6815 SECCLASS 255 NOMONSYSPLEX IPADDR
          2001:DB8:10::11:1:5
INTERFace MPC6245 DEFINE MPCPTP6 TRLE MPC4245 INTFID ::12:4:5
          SOURCEVIPAINTerface VIPA6815 SECCLASS 255 NOMONSYSPLEX IPADDR
          2001:DB8:10::12:4:5
INTERFace MPC6185 DEFINE MPCPTP6 TRLE MPC4185 INTFID ::11:8:5
          SOURCEVIPAINTerface VIPA6815 SECCLASS 255 NOMONSYSPLEX IPADDR
          2001:DB8:10::11:8:5
INTERFace QDIO6105 DEFINE IPAQENET6 CHPIDTYPE OSD PORTNAME QDIO4105 PRIROUTER
          DUPADDRDET 1 SOURCEVIPAINTerface VIPA6815 NODYNVLANREG
          READSTORAGE GLOBAL INBPERF BALANCED INTFID ::16:1:5 SECCLASS
          255 NOMONSYSPLEX TEMPPREFIX ALL IPADDR 2001:DB8:172::16:1:5 NOISOLATE NOOLM
IPCONFig ARPTO 1200 DATAGRamfwd NOFWDMULTipath DEVRETRYDURation 4
         IPSECurity NOSOURCEVIPA NOTCPSTACKSOURCEVIPA NOSYSPLEXRouting
         REASSEMBLytimeout 60 TTL 64 NOPATHMTUDISCovery NOMULTIPATH
         NODYNAMICXCF NOIQDIORouting FORMAT LONG NOQDIOACCELerator
         NOSEGMENTATIONOFFLOAD CHECKSUMOFFLOAD
IPCONFIG6 DATAGRamfwd NOFWDMULTipath IPSECurity NOSOURCEVIPA
          NOTCPSTACKSOURCEVIPA NOMULTIPATH HOPLimit 255 ICMPErrorlimit
          3 NOIGNOREROUTERHOPLIMIT NODYNAMICXCF NOTEMPADDRS
          NOSEGMENTATIONOFFLOAD CHECKSUMOFFLOAD
IPSEC DVIPsec LOGENable LOGImplicit
  IPSECRule * * NOLOG PROTOcol OSPF TYPE * ROUTING LOCAL SECCLASS 0
  IPSECRule * * LOG PROTOcol * ROUTING LOCAL SECCLASS 0
  IPSEC6Rule * * NOLOG PROTOcol * ROUTING LOCAL SECCLASS 0
ENDIPSEC
ITRACE OFF AUTODAEMON
ITRACE OFF COMMAND
ITRACE OFF CONFig
ITRACE OFF SUBAGENt
PORT 3020 TCP DCICSTS DELAYAcks
PORT 3019 TCP DCICSTS DELAYAcks
PORT 3018 TCP DCICSTS DELAYAcks
PORT 3017 TCP DCICSTS DELAYAcks
PORT 3016 TCP DCICSTS DELAYAcks
PORT 3015 TCP DCICSTS DELAYAcks
PORT 3014 TCP DCICSTS DELAYAcks
PORT 3013 TCP DCICSTS DELAYAcks
```

```
PORT 3012 TCP DCICSTS DELAYAcks
PORT 3011 TCP DCICSTS DELAYAcks
PORT 3010 TCP DCICSTS DELAYAcks
PORT 3009 TCP DCICSTS DELAYAcks
PORT 3008 TCP DCICSTS DELAYAcks
PORT 3007 TCP DCICSTS DELAYAcks
PORT 3006 TCP DCICSTS DELAYAcks
PORT 3005 TCP DCICSTS DELAYAcks
PORT 3004 TCP DCICSTS DELAYAcks
PORT 3003 TCP DCICSTS DELAYAcks
PORT 3002 TCP DCICSTS DELAYAcks
PORT 3001 TCP DCICSTS DELAYAcks
PORT 3000 TCP DCICSTS DELAYAcks
PORT 520 UDP * DELAYAcks
PORT 161 UDP OSNMPD DELAYAcks
PORT 53 UDP OMVS DELAYAcks
PORT 53 TCP OMVS DELAYAcks
PORT 23 TCP TCPCS5 DELAYAcks
PORT 21 TCP OMVS DELAYAcks
PORT 19 TCP MISCSERV DELAYAcks
PORT 19 UDP MISCSERV DELAYAcks
PORT 9 TCP MISCSERV DELAYAcks
PORT 9 UDP MISCSERV DELAYAcks
PORT 7 TCP MISCSERV DELAYAcks
PORT 7 UDP MISCSERV DELAYAcks
SACONFig COMMUNity public AGENT 161 SACACHETime 30 ENABLed SETSDISAbled
         OSADISabled
SMFCONFIG TYPE118 NOTCPINIT NOTCPTERM NOFTPCLIENT NOTN3270CLIENT
          NOTCPIPStatistics TYPE119 NOTCPINIT NOTCPTERM NOFTPCLIENT
          NOTN3270CLIENT NOTCPIPStatistics NOIFStatistics
          NOPORTStatistics NOTCPSTACK NOUDPTerm NOIPSECURITY
          PROFILE DVIPA
SOMAXCONN 10
START IUTSAMEH
START IUTSAMEH6
STOP  MPC4115
STOP  MPC6115
STOP  MPC4145
STOP  MPC4185
STOP  MPC6185
STOP  QDIO4105I
STOP  I4HIPERSOCKA
STOP  QDIO6105
STOP  MPC6245
TCPCONFIG INTerval 120 DELAYAcks RESTRICTLowports TCPRCVBufrsize 16384
          TCPSENDBfrsize 16384 TCPMAXRCVBufrsize 262144 FINWAIT2TIME
          600 SENDGarbage FALSE TCPTIMESTAMP NOTTLS SELECTIVEACK
          EPHEMERALPORTS 01024 - 65535
UDPCONFIG RESTRICTLowports UDPCHKsum UDPRCVBufrsize 65535
          UDPSENDBfrsize 65535 UDPQueuelimit EPHEMERALPORTS 01024 - 65535
BEGINRoutes
 ROUTE 10.81.4.0 255.255.255.0 10.11.5.4 MPC4145L MTU DEFAULTSIZE
       MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks NOREPLaceable
 ROUTE 10.81.2.0 255.255.255.0 10.11.5.4 MPC4145L MTU DEFAULTSIZE
       MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks NOREPLaceable
 ROUTE 10.51.0.4 HOST = TOVTAM MTU DEFAULTSIZE MAXImumretransmittime
       120 MINImumretransmittime 0.5 ROUNDTRIPGain 0.125 VARIANCEGain
       0.25 VARIANCEMultiplier 2 DELAYAcks NOREPLaceable
 ROUTE 10.11.5.1 HOST = MPC4115L MTU 4096 MAXImumretransmittime 120
       MINImumretransmittime 0.5 ROUNDTRIPGain 0.125 VARIANCEGain 0.25
       VARIANCEMultiplier 2 DELAYAcks REPLaceable
 ROUTE 10.11.5.4 HOST = MPC4145L MTU 4096 MAXImumretransmittime 120
       MINImumretransmittime 0.5 ROUNDTRIPGain 0.125 VARIANCEGain 0.25
```

```
       VARIANCEMultiplier 2 DELAYAcks REPLaceable
 ROUTE 10.11.5.8 HOST = MPC4185L MTU 4096 MAXImumretransmittime 120
       MINImumretransmittime 0.5 ROUNDTRIPGain 0.125 VARIANCEGain 0.25
       VARIANCEMultiplier 2 DELAYAcks REPLaceable
 ROUTE 10.81.8.0 255.255.255.0 10.11.5.8 MPC4185L MTU DEFAULTSIZE
       MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks NOREPLaceable
 ROUTE 172.16.0.0 255.255.0.0 = QDIO4105I MTU 1492
       MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks REPLaceable
 ROUTE 2001:DB8:10::81:4:4/128 = MPC6245 MTU DEFAULTSIZE
       MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks NOREPLaceable
 ROUTE 2001:DB8:10::81:8:0/112 2001:DB8:10::11:5:8 MPC6185 MTU
       DEFAULTSIZE MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks NOREPLaceable
 ROUTE 2001:DB8:10::51:0:4/128 = IUTSAMEH6 MTU DEFAULTSIZE
       MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks NOREPLaceable
 ROUTE 2001:DB8:10::11:5:1/128 = MPC6115 MTU 4096 MAXImumretransmittime
       120 MINImumretransmittime 0.5 ROUNDTRIPGain 0.125 VARIANCEGain
       0.25 VARIANCEMultiplier 2 DELAYAcks REPLaceable
 ROUTE 2001:DB8:10::12:5:4/128 = MPC6245 MTU 4096 MAXImumretransmittime
       120 MINImumretransmittime 0.5 ROUNDTRIPGain 0.125 VARIANCEGain
       0.25 VARIANCEMultiplier 2 DELAYAcks REPLaceable
 ROUTE 2001:DB8:10::11:5:8/128 = MPC6185 MTU 4096 MAXImumretransmittime
       120 MINImumretransmittime 0.5 ROUNDTRIPGain 0.125 VARIANCEGain
       0.25 VARIANCEMultiplier 2 DELAYAcks REPLaceable
 ROUTE 2001:DB8:10::81:2:0/112 2001:DB8:10::12:5:4 MPC6245 MTU
       DEFAULTSIZE MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks NOREPLaceable
 ROUTE 2001:DB8:172::16:0:0/96 = QDIO6105 MTU 1492
       MAXImumretransmittime 120 MINImumretransmittime 0.5
       ROUNDTRIPGain 0.125 VARIANCEGain 0.25 VARIANCEMultiplier 2
       DELAYAcks REPLaceable
ENDRoutes
Analysis of Tcp/Ip for TCPCS5 completed
```

TCPIPCS PROTOCOL

Use this subcommand to display information from TCP, UDP, and RAW protocol control blocks.

Syntax

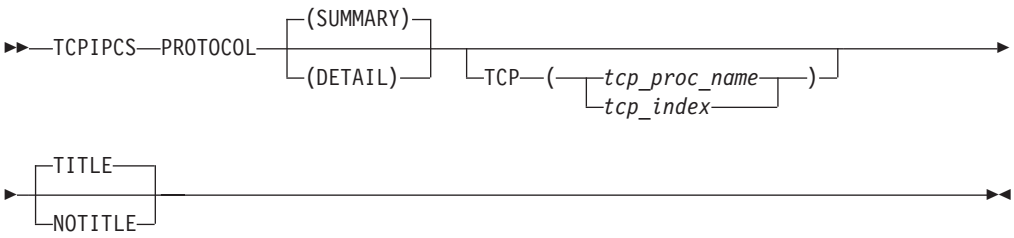

Parameters

SUMMARY
Formats the MTCB, MUDP, and MRCB contents. Lists all the TCBs, UCBs, and RCBs in separate cross-referenced tables. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of the TCBs, UCBs, and RCBs.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS PROTOCOL subcommand

The following is sample output of the TCPIPCS PROTOCOL subcommand:

```
TCPIPCS PROTOCOL
 Dataset: IPCS.MV21381.DUMPA
 Title:   SLIP DUMP ID=TC


 The address of the TSAB is: 13391BC0

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 13391C00  1 TCPSVT    V2R10   1323B000 1323B0C8 07DE 04041405  Active
 13391C80  2 TCPSVT2   V2R10   00000000 00000000 07E8 00000000  Down Stopping
 13391D00  3 TCPSVT1   V2R10   12FC3000 12FC30C8 0080 94FF755F  Active
 13391D80  4 TCPSVT3   V2R10   00000000 00000000 0059 00000000  Down Stopping

    4 defined TCP/IP(s) were found
    2 active  TCP/IP(s) were found

    4 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPSVT.  Index: 1


 TCPIP Raw Control Block Analysis
 Master Raw Control Block (MRCB)
  MRAWCB: 7F75B048
  +0000  RMRCBEYE. MRCB      MRCMUTEX. 00000000  00000000  00000000
  D7D60501            RSTKDOWN. 00
  +0021  RSTKLNKD. 01        RDRVSTAT. 01        RSBCAST.. 00000000
  RSDNTRTE. 00000000  RSRCVBUF. 0000FFFF
  +0030  RSSNDBUF. 0000FFFF  RDIPTOS.. 00        RDIPTTL.. 00
  RIPWRQ@.. 7F61D3E8  RIPRDQ@.. 7F61D3A8
  +0040  RHASH@... 7F75B08C
....

Raw Hash Table Entries

ID  First    Last
9   7F5513C8 7F5513C8
15  7F712088 7F712088


RCB      ResrcID  ResrcNm  TpiState   DestAddr            ProtocolId
7F5513C8 00000062 OMPROUTE WLOIDLE    129.11.208.108      89
7F712088 00000008 TCPSVT   WLOIDLE    0.0.0.0             255
```

```
2 RCB(s) FOUND
2 RCB(s) FORMATTED


-------------------------------------------------------------------------------
TCP/IP Analysis
TCPIP Main TCP Control Block (MTCB)
 MTCB: 1338E350
 +0000  M_MAIN_EYE......... TCP MAIN
 +0008  M_TCP_LWRITE_Q..... 7F781868
 +000C  M_TCP_LREAD_Q...... 7F781828
 +0014  M_TCP_DRIVER_STATE. 01
 +0018  MTCPMTX............ 00000000  00000000  00000000  D7D60601
 +0028  MTCPAQMX........... 00000000  00000000  00000000  D7D60604
 +0038  MTCB_LIST_LOCK..... 00000000  00000000  00000000  D7D60604
 +0048  M_PORT_CEILING..... 00000FFF
 +004C  M_TPI_SEQ#......... 0001C62B
 +0050  M_PORT_ARRAY....... 7F712FC8
 +0054  M_LAST_PORT_NUM.... 00000445
.....

 TCB      ResrcID  ResrcNm  TcpState    TpiState Flag1234 UseCount IPAddr
     Port LuName   ApplName UserID
 7F607108 00000002 TCPSVT   Closed      WLOUNBND 00040000 00000001 0.0.0.0
        0
 7F60A908 000083D7 FTPUNIX1 Listening   WLOIDLE  00200080 00000001 0.0.0.0
        0
 7F608D08 00000013 TCPSVT   Listening   WLOIDLE  00000080 00000001 0.0.0.0
        0
 7F617508 0000019B CICSRU   Listening   WLOIDLE  08200080 00000001 0.0.0.0
        0
 7F615108 00000144 INETD5   Listening   WLOIDLE  00200080 00000001 0.0.0.0
        0
...
 7F610108 0000878F NAMED4   TimeWait    WLOWIORL 80800C00 00000002 198.11.22.103
      53
 7F60C508 0000005C DHCP1    Established WLOXFER  01800000 00000001 198.11.25.104
    6000
 7F609D08 00000049 MISCSRV  Listening   WLOIDLE  00200000 00000001 0.0.0.0
       0
 7F608908 00000012 TCPSVT   Listening   WLOIDLE  00000080 00000001 0.0.0.0
       0
 7F60E108 00000063 TCPSVT   Established WLOXFER  80800000 00000001 127.0.0.1
    1030
 30 TCB(s) FOUND
 30 TCB(s) FORMATTED
 -------------------------------------------------------------------------------
 User Datagram Protocol Control Block Summary
  MUCB: 7F7812A8
  +0000  UMUCBEYE. MUCB      USTKDOWN. 00        USTKLNKD. 01
  UAPAR.... 00        UDRVSTAT. 00
  +0008  UOPENPRT. 00000000  UFREEPRT. 0408      MCBMUTEX. 00000000
  00000000  00000000  D7D60402
  +0020  UDPCFG... 00000001  0000FFFF  0000FFFF  00000001  80000000
  00000000
  +0038  UDPCFG2.. 00000001  0000FFFF  0000FFFF  00000001  80000000
  00000000
  +0050  UDPMIB... 00001D1F  0000531F  00000000  0000166B
  USBCAST.. 00000000  USLPBACK. 00000000
  +0068  USDNTRTE. 00000000  USRCVBUF. 0000FFFF  USSNDBUF. 0000FFFF
  UFGPRC... 00        USERIALV. 0000065F
  +007C  USERIAL1. 0000065F  ULASTADR. 810B2068  ULASTPRT. 0043
  ULASTUCB. 7F5FD508  USERIAL2. 0000065F
    ...

 UCB      ResrcID  ResrcNm  TpiState IPAddr          Port
 7F5F6108 00000004 TCPSVT   WLOUNBND 0.0.0.0
```

```
 7F5FCD08 00000086 OSNMPD   WLOIDLE  127.0.0.1        161
 7F5FD508 0000005E DHCP1    WLOIDLE  129.11.32.1       67
 7F5FCF08 00000055 DHCP1    WLOIDLE  198.11.25.104   1027
 7F5FD308 0000005B NAMED    WLOIDLE  129.11.176.87     53
 7F5FD108 00000059 DHCP3    WLOIDLE  0.0.0.0         6001
 7F5FC908 00000048 MISCSRV  WLOIDLE  0.0.0.0           19
 7F5FC708 00000047 MISCSRV  WLOIDLE  0.0.0.0           19
....
 7F5F6B08 00000017 MISCSRV  WLOIDLE  0.0.0.0            7
 7F5F6908 00000014 PORTMAP  WLOIDLE  0.0.0.0          111
 56 UCB(s) FOUND
 56 UCB(s) FORMATTED

 Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS RAW

Use this subcommand to display the Master Raw Control Block (MRCB) and any Raw protocol Control Blocks (RCBs) defined in the MRCB hash table.

Syntax

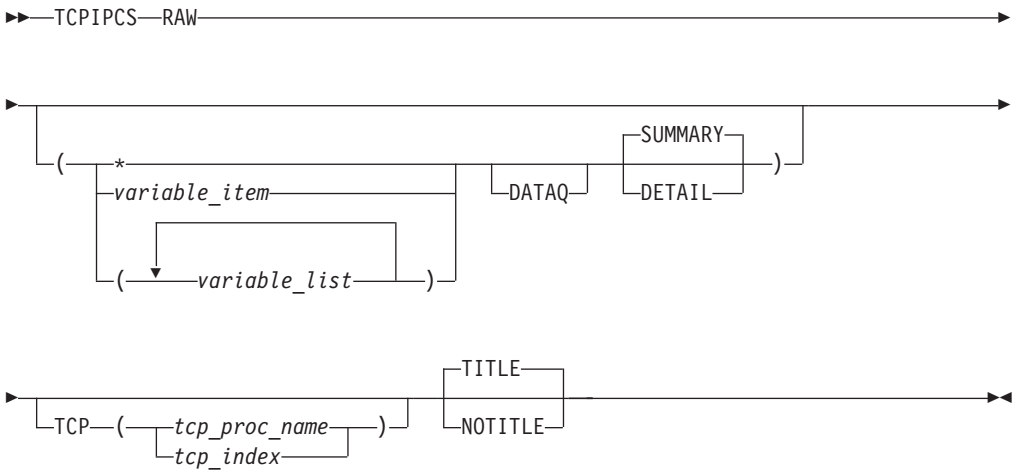

Parameters

If no parameters are specified, all raw connections are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

jobname
Displays only the RCB for this job name. The job name can be a TCP/IP application name or a stack name, and it must contain from 1-8 characters.

RCB_address
Displays only the RCB with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address.

If an address begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

connection_id
Displays the RCB with this connection ID. A connection ID is specified as 1 - 8 hexadecimal digits.

In addition to the variable parameters described above, you can specify the following keyword parameters:

DATAQ
Formats RCBs which have data queued on the RECEIVE queue.

SUMMARY
Formats the MRCB contents and lists all the RCBs in one cross-reference table. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of the RCBs.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS RAW subcommand

The following is sample output of the TCPIPCS RAW subcommand:

```
TCPIPCS RAW
Dataset: IPCS.R8A0723.RASDUMP
Title:   EZRPE005
The address of the TSAB is: 098221F0
Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status
09822230  1 TCPCS     V1R5    08E85000 08E850C8 001E 9FFF7E7F  Active
098222B0  2 TCPCS2    V1R5    08937000 089370C8 01F6 9FFF7E7F  Active
   2 defined TCP/IP(s) were found
   2 active  TCP/IP(s) were found
   2 TCP/IP(s) for CS     V1R5  found

Analysis of Tcp/Ip for TCPCS.  Index: 1
TCPIP Raw Control Block Analysis
Master Raw Control Block (MRCB)
 MRAWCB: 7F407208
 +0000  RMRCBEYE. MRCB      RSTKLNKD. 01        RDRVSTAT. 01
 +000C  R6STKLNKD.                    01
 +000D  R6DRVSTAT.                    01
 +0010  MRCMUTEX. 00000000  00000000  00000000  D7D60501
        RSBCAST.. 00000000  RSDNTRTE. 00000000  RSRC
 +002C  RSSNDBUF. 0000FFFF  RDIPTOS.. 00        RDIPTTL.. 00
        RIPWRQ@.. 7F621DA8  RIPRDQ@.. 7F621D68  RHAS
 +0040  RIP6WRQ@. 7F686468  RIP6RDQ@. 7F686428  R6HASH@.. 7F407374
 +004C  R6DFFLTR. 7F781FFF  FFFFFFFF  FFFFFFFF  FFFFFFFF  003FFFFF
        FFFFFFFF  FFFFFFFF  FFFFFFFF

IPv4 Raw Hash Table Entries
ID  First    Last
0   7F52C390 7F52C390
15  7F52C110 7F52C110

IPv6 Raw Hash Table Entries
ID  First    Last
0   7F2073BC 7F2073BC
```

```
IPv4 RAW Connections
RCB      ResrcID  ResrcNm  TpiState   ProtocolId  DestAddr
7F52C388 00000006 TCPCS    WLOIDLE    0           0.0.0.0
7F52C108 00000008 TCPCS    WLOIDLE    255         0.0.0.0

IPv6 RAW Connections
RCB      ResrcID  ResrcNm  TpiState   ProtocolId  DestAddr
7F207208 0000000E TCPCS    WLOIDLE    0           ::0

3 RCB(s) FOUND
3 RCB(s) FORMATTED

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS ROUTE

Use this subcommand to display the routing control blocks. Each routing table entry is formatted to display the:

- Route control block address
- Device name
- Type
- Protocol
- Destination IP address
- Gateway IP address
- Physical interface control block address

Syntax

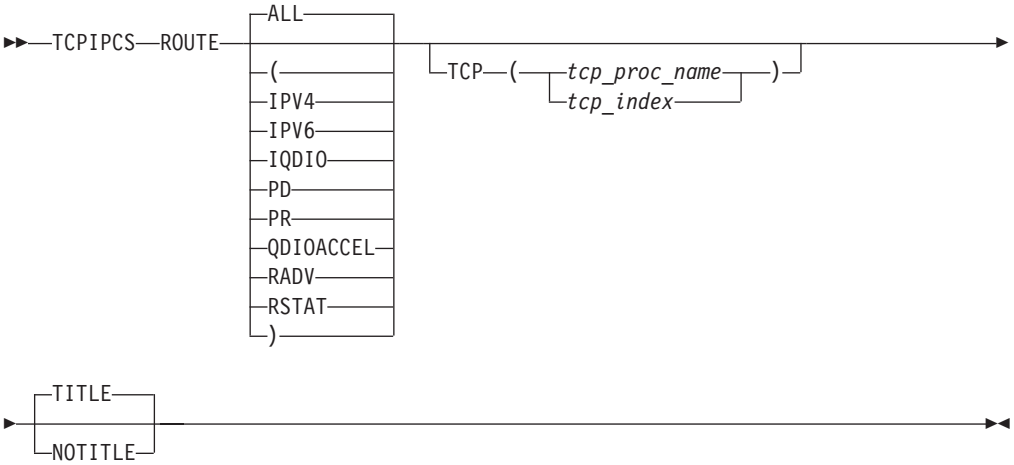

Parameters

ALL
Display structure of all route table information (including all active and to-be-deleted policy-based routing tables). ALL is the default.

IPV4
All IPv4 search tree and update tree routes in the main route table.

IPV6
All IPv6 search tree and update tree routes in the main route table.

IQDIO
All QDIO Accelerator and HiperSockets™ Accelerator search tree and update tree routes.

PD All policy-based routing tables that are marked for deletion.

PR All search tree and update tree routes for all policy-based routing tables. Also, list configured routes that use interfaces that are not defined in the stack and list dynamic routing parameters for all policy-based routing tables.

QDIOACCEL
All QDIO Accelerator and HiperSockets Accelerator search tree and update tree routes.

RADV
All IPv6 routes added based on information that was received in router advertisement messages are displayed without regard to whether they are currently being used in the active routing table.

RSTAT
All defined replaceable static routes are displayed without regard to whether they are currently being used in the active routing table.

TCP, TITLE, NOTITLE
For a description of these parameters, see "Parameters" on page 177.

Rule: If you specify multiple keywords from the set {ALL, IPV4, IPV6, IQDIO, PD, PR, QDIOACCEL, RADV, RSTAT}, all of the keywords that you specify are used.

Sample output of the TCPIPCS ROUTE subcommand

The following is sample output of the TCPIPCS ROUTE subcommand:

```
TCPIPCS ROUTE
Dataset: IPCS.P414001.PUBDUMPA
Title:   DUMP OF TCP STACKS


The address of the TSAB is: 3A467000

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts          Status

3A467040  1 TCPSVT    V1R9    3A449000 3A4490C8 006A 97BF749F C0000000  Active
3A4670C0  2 TCPSVT1   V1R9    397CD000 397CD0C8 007B 97BF749F C0000000  Active
3A467140  3 TCPSVT2   V1R9    38669000 386690C8 007C 97BF749F C0000000  Active

   3 defined TCP/IP(s) were found
   3 active  TCP/IP(s) were found

   3 TCP/IP(s) for CS     V1R9  found

================================================================================

Analysis of Tcp/Ip for TCPSVT.  Index: 1

TCPIP Route Analysis

================================================================================
Replaceable Static Routes Configured

  IPv4 Route Table Name=EZBMAIN  ID=0 Ref=0 Flags=08000000

    Rtioctl@  LinkName          IP Addresses
    --------- ----------------  ------------------
    7F453150  LOGETH2           Destination: 174.33.84.237
                                Gateway    : 0.0.0.0
    7F452BB0  LOGETH2           Destination: 197.0.0.0
                                Gateway    : 174.33.84.237
```

```
  IPv6 Route Table Name=EZBMAIN  ID=0 Ref=0 Flags=08080000

    Rtioclt6@ InterfaceName     IP Addresses
    --------- ----------------  -----------------
    7F452750  LV6OGETH2         Destination: 2000:176:11:48::237
                                Gateway    : ::0

================================================================================
IPv6 Router Advertisement Routes

  IPv6 Route Table Name=EZBMAIN  ID=0 Ref=0 Flags=08080000

================================================================================
Routes in IPv4 Route Tables

  IPv4 Route Table Name=EZBMAIN  ID=0 Ref=0 Flags=08000000

    Routes in Search Table

      Rte@     LinkName         Type/State      Protocol        Pif@     IP Addresses
      -------- ---------------- --------------- --------------- -------- -----------------
      7DCABB50 LOGETH2          Host            OSPF            7F850490 Destination: 202.77.232.1
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 174.33.84.237
      7DFC25D0 LOGETHB          Host            OSPF            7F850090 Destination: 202.77.232.1
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 174.33.84.237
      7DCABE90 LOGETH2          Host            OSPF            7F850490 Destination: 202.77.230.1
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 174.33.84.237
      7DCABCF0 LOGETHB          Host            OSPF            7F850090 Destination: 202.77.230.1
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 174.33.84.237
      7F453350 LOGETH2          Host            Configuration   7F850490 Destination: 174.33.84.237
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 0.0.0.0
      7F452E70 LOGETH2          Subnetwork      Configuration   7F850490 Destination: 197.0.0.0
                                Active                                   Subnet Mask: 255.0.0.0
                                                                         Gateway    : 174.33.84.237

    Routes in Update Table

      Rte@     LinkName         Type/State      Protocol        Pif@     IP Addresses
      -------- ---------------- --------------- --------------- -------- -----------------
      7DCABB50 LOGETH2          Host            OSPF            7F850490 Destination: 202.77.232.1
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 174.33.84.237
      7DFC25D0 LOGETHB          Host            OSPF            7F850090 Destination: 202.77.232.1
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 174.33.84.237
      7DCABE90 LOGETH2          Host            OSPF            7F850490 Destination: 202.77.230.1
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 174.33.84.237
      7F453350 LOGETH2          Host            Configuration   7F850490 Destination: 174.33.84.237
                                Active                                   Subnet Mask: 255.255.255.255
                                                                         Gateway    : 0.0.0.0
      7F452E70 LOGETH2          Subnetwork      Configuration   7F850490 Destination: 197.0.0.0
                                Active                                   Subnet Mask: 255.0.0.0
                                                                         Gateway    : 174.33.84.237

================================================================================
Routes in IPv6 Route Tables

  IPv6 Route Table Name=EZBMAIN  ID=0 Ref=0 Flags=08080000

    Routes in Search Table

      Rte@     InterfaceName    Type/State      Protocol        Pif@     IP Addresses
      -------- ---------------- --------------- --------------- -------- -----------------
```

```
   7DCA7730 LV6OGETH2       Host           OSPF           7F454650 Destination: fec0:197:11:104::20
                            Active                                 Prefix     : 124
                                                                   Gateway    : fe80::11:176:50:104
   7EBFB450 LV6IUTIQD00     Host           Configuration  7F455250 Destination: fe80::2440:ff:f10c:2
                            Active                                 Prefix     : 128
                                                                   Gateway    : ::0
   7EBFB2B0 LV6IUTIQD02     Host           Configuration  7F4567F0 Destination: fe80::2440:2ff:f10c:42
                            Active                                 Prefix     : 128
                                                                   Gateway    : ::0
   7EB95ED0 EZ6SAMEMVS      Host           Configuration  7EB91010 Destination: fe80::11:199:80:104
                            Active                                 Prefix     : 128
                                                                   Gateway    : ::0
   7F4528D0 LV6OGETH2       Host           Configuration  7F454650 Destination: 2000:176:11:48::237
                            Active                                 Prefix     : 128
                                                                   Gateway    : ::0

 Routes in Update Table

   Rte@     InterfaceName    Type/State     Protocol       Pif@     IP Addresses
   -------- ---------------- -------------- -------------- -------- ------------------
   7DCA7730 LV6OGETH2       Host           OSPF           7F454650 Destination: fec0:197:11:104::20
                            Active                                 Prefix     : 124
                                                                   Gateway    : fe80::11:176:50:104
   7EBFB450 LV6IUTIQD00     Host           Configuration  7F455250 Destination: fe80::2440:ff:f10c:2
                            Active                                 Prefix     : 128
                                                                   Gateway    : ::0
   7EBFB2B0 LV6IUTIQD02     Host           Configuration  7F4567F0 Destination: fe80::2440:2ff:f10c:42
                            Active                                 Prefix     : 128
                                                                   Gateway    : ::0
   7F4528D0 LV6OGETH2       Host           Configuration  7F454650 Destination: 2000:176:11:48::237
                            Active                                 Prefix     : 128
                                                                   Gateway    : ::0

Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS SOCKET

Use this subcommand to display information from TCP/IP socket control blocks.

Syntax

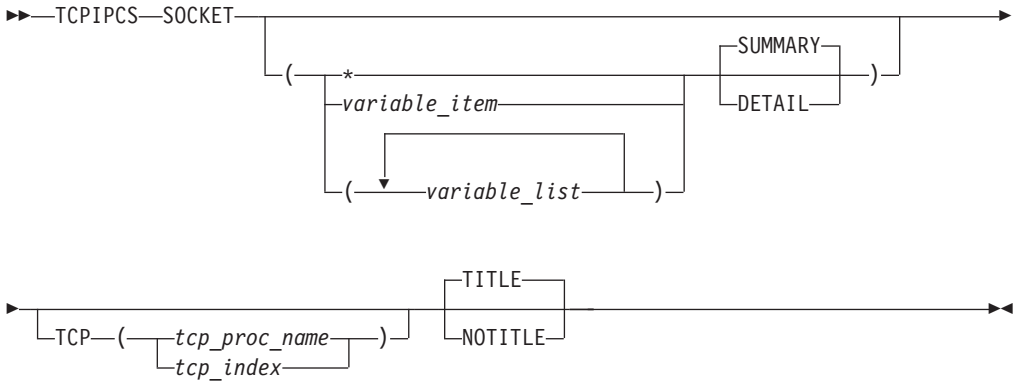

Parameters

If no parameters are specified, all sockets are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1–32 of the following variable parameters, each separated by a blank space, within parentheses:

SCB_address
Displays only the socket control block (SCB) with this address. An address is specified as 1-16 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

connection_id
Displays the SCB with this connection ID. A connection ID is specified as 1-8 hexadecimal digits.

In addition to the variable parameters described above, the following keyword parameters can be specified:

SUMMARY
Summarizes the sockets. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of the SCBs.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS SOCKET subcommand

The following is sample output of the TCPIPCS SOCKET subcommand:

```
TCPIPCS SOCKET
      Dataset: IPCS.MV21381.DUMPA
      Title:   SLIP DUMP ID=TC


      The address of the TSAB is: 13391BC0

      Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status


      13391C00  1 TCPSVT    V2R10   1323B000 1323B0C8 07DE 04041405  Active
      13391C80  2 TCPSVT2   V2R10   00000000 00000000 07E8 00000000  Down Stopping
      13391D00  3 TCPSVT1   V2R10   12FC3000 12FC30C8 0080 94FF755F  Active
      13391D80  4 TCPSVT3   V2R10   00000000 00000000 0059 00000000  Down Stopping


         4 defined TCP/IP(s) were found
         2 active  TCP/IP(s) were found


         4 TCP/IP(s) for CS     V2R10 found

      ================================================================================

      Analysis of Tcp/Ip for TCPSVT.  Index: 1

      TCPIP Socket Analysis

      SCB              CID      Protocol    SockOpts ScbFlags ResrcNm
      0000000112D42120 00000008 RAW         00000000 00280000 TCPSVT
```

```
0000000112D422A0 0000000B UDP         00000000 00280000 TCPSVT
0000000112D42420 0000000C TCP         00020000 C0280000 TCPSVT
0000000112D425A0 0000000E UDP         00000000 00280000 TCPSVT
0000000112D42720 0000000F TCP         00000000 B0280000 TCPSVT
0000000112D428A0 00000010 TCP         00020000 90280000 TCPSVT
0000000112D42A20 00000067 TCP         08000000 90280000 OMPROUTE
0000000112D42BA0 00000012 TCP         00400000 C0000000 TCPSVT
0000000112D42D20 00000013 TCP         00400000 C0000000 TCPSVT
0000000112D42EA0 00000014 UDP         00000000 80280000 PORTMAP
0000000112D43020 00000015 TCP         00000000 C0280000 PORTMAP
 ...
0000000112D44FA0 00000058 TCP         00000000 C0000000 DHCP3
0000000112D45120 00000059 UDP         00000000 80280000 DHCP3
0000000112D452A0 0000005A TCP         00400000 C0280000 NAMED
0000000112D45420 0000005B UDP         00400000 80280000 NAMED

79 Socket control blocks were found
79 Socket control blocks were formatted


Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS STATE

Use this subcommand to provide an overall view of TCP/IP. The following information is displayed:

- Major control block addresses
- Subtasks
- Storage usage
- Dispatchable units
- Trace
- Configuration

Syntax

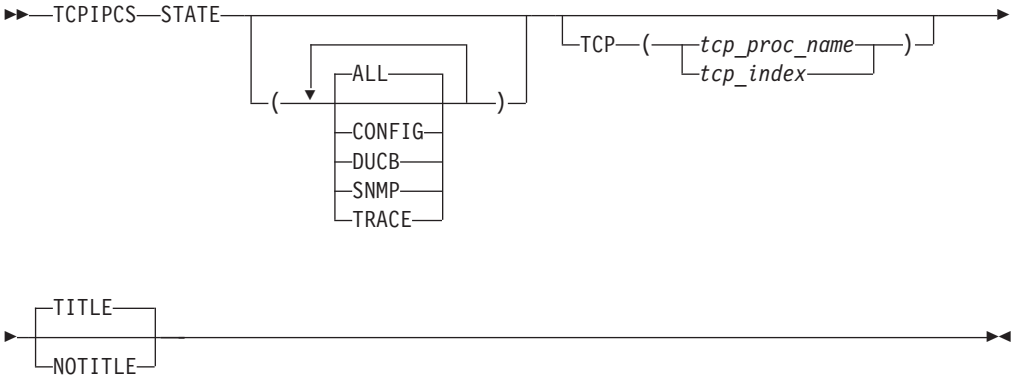

Parameters

ALL
Display all state information. ALL is the default.

CONFIG
Display only configuration state information.

DUCB
Display only DUCB state information.

SNMP
Display only SNMP and CONFIG information. (SNMP information makes sense only in the context of the configuration, so the configuration information is also displayed.)

TRACE
Display only trace state information.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {ALL, CONFIG, DUCB, SNMP, TRACE}, all of the keywords that you specify are used.

Sample output of the TCPIPCS STATE subcommand

The following is sample output of the TCPIPCS STATE subcommand:

```
--------------------------------------------------------------------------------

Dataset: IPCS.X370812.TRAK0006.STEP2B.DUMP
 Title:   TRAK0006 STEP1


 The address of the TSAB is: 15137000

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts          Status

 15137040  1 TCPCS     V1R8    14F5A000 14F5A0C8 0030 9FFF767F 00000000  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V1R8  found

 ================================================================================

 Analysis of Tcp/Ip for TCPCS.  Index: 1

 TCPIP State

 TCPIP Status:
  Procedure:     TCPCS
  Version:       V1R8
  Status:        Active
  Asid:          0030
  Started:       2005/08/24 16:21:15
  Ended:         2005/08/24 16:38:29
  Active:        00:17:13.628105 hours

 Major Control Blocks
  TSEB:          15137040       TSDB:          14F5A000
  TSDX:          14F5A0C8       TCA:           15A59418
  ITCVT:         14F5A398       ITSTOR:        14F5A608
  DUAF:          14F56018       MRCB:          7F41D0F0
  MTCB:          15C6D2F8       MUCB:          7F3F30F0
  IPMAIN:        14E67398       Streams_root:  7F5DAD10
  TosMains:      15AA3C48       MIB2:          15A59258
  CdCb:          15D98260       User:          15DAC000
  Conf:          15A56A98       Stks:          15D981A0
  IPMAIN6:       15AA3848
 ================================================================================
 TCPIP Subtasks
 Task      Tcb       FirstRB   EotECB    StopEcb   CmpCode   RsnCode   RTWA
 EZBTCPIP  008FF2A0  008FF218  808FD2D0            00000000  00000000  00000000
 EPWPITSK  008E4E88  008DC088  00000000            00000000  00000000  00000000
 EZBITTUB  008E4CF0  008EB340  00000000  808DC198  00000000  00000000  00000000
```

```
EZACDMSM  008DE088  008FF4A8  00000000  808FF4A8  00000000  00000000  00000000
EZBIPSUB  008E4B58  008DC110  00000000  808DC110  00000000  00000000  00000000
EZBIEOER  008E47D8  008E4710  008FF218  808E4710  940C4000  00000004  00000000
EZBTLMST  008E4578  008E44D0  008FF218            00000000  00000000  00000000
EZACFMMN  008E4240  008E41B8  808FF218  808E41B8  00000000  00000000  00000000
EZBTZMST  008EB470  008E4038  008FF218            00000000  00000000  00000000
EZBTTSSL  008DED90  008DED08  808E4038            00000000  00000000  00000000
EZBTMCTL  008DEA78  008DEF68  808E4038            00000000  00000000  00000000
EZACFALG  008DE8E0  008DE858  808FF218            00000000  00000000  00000000
EZASASUB  008DCE88  008DEC70  808FF218  808DEC70  00000000  00000000  00000000
================================================================================

Storage Cache Information
  Total CSA Allocated:                7,705,448
  Tcp/ip CSA Limit:                     2,147M
  Total CSA Elements:                       59
  Cache Delay:                             210 seconds
  Scan Delay:                               75 seconds
  Total cache allocated:                91,688
  Total cache elements:                      8
  Total freed elements:                      0
  Last cache scan time:     2005/08/24 20:37:12

CSM Status
  ECSA Storage:           OK
  Data Space Storage:     OK
  Fixed Storage:          OK
Alet: 01FF0014               Dspname: CSM64001
Alet: 00000000               Dspname: ........

================================================================================
Dispatchable Unit Status
  DUCB Initializations:                11,741
  DUCB Expansions:                        769
  Percent DUCB expansions:                  6 %
  Last DUCB scan time:     2005/08/24 20:34:24

   1 DUAT control block(s) were found in the DUAF at 14F56018
 124 Dispatchable units were found.
  No DUs indicate abend.
================================================================================

CTrace Status:
   Member Name  : CTIEZBN0
   Buffer Size  : 4,194,304
   Options      : Init Opcmds Opmsgs Socket AFP XCF Access PFS
                  API Engine Queue RAW UDP TCP ICMP ARP ND CLAW
                  LCS Internet Message WorkUnit Config SNMP
                  IOCTL FireWall VtamData TelnVtam Telnet Vtam
   Asid List    : ()
   JobNameList  : ()
   PortList     : ()
   IpAddrList   : ()
   Xwriter      : Disconnected
   Dwriter      : Disconnected
   Trace Count  : 25,553
   Lost Count   : 2
   Lost Time    : 2005/08/24 20:21:16
   Wrap Count   : 1
   Wrap Time    : 2005/08/24 20:36:05
================================================================================

Device Interface: 7F5DC410
  Device: LOOPBACK         Devtype: LOOPBACK         State: Active
  Address: **** ****

Physical Interface: 7F5DA230
  Name: LOOPBACK          Protocol: LOOPBACK         State: Active
```

```
    NetNum: 0  QueSize: 0  Bytein: 13,554     Byteout: 13,554
     Index: 2
  Bsd Routing Parameters:
   MtuSize: 0                   Metric: 0
   SubnetMask: 0.0.0.0          DestAddr: 0.0.0.0
  SNMP Input Counters:
         Octets:               13,554    Unicast:                  214
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Unkn Type:                    0
      Broadcast:                    0  Multicast:                    0
  SNMP Output Counters:
         Octets:               13,554    Unicast:                  214
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Queue Len:                    0
      Broadcast:                    0  Multicast:                    0

  IPv4 Search Patricia tree     Address: 7F55CF10
    Search Ptree Reader Count: 0

     Route: 7F5DC270
      Name: LOOPBACK          Type: Host             State: Active
      Subnet Mask: 255.255.255.255  Addr: 127.0.0.1
      Protocol   : Configuration    Gate: 0.0.0.0
      Mtu Size: 65535         Ref Cnt: 4          Tos: 0
      Metric1:  0             Metric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:19
  IPv6 Search Patricia tree     Address: 7F55C370
    Search Ptree Reader Count: 0



Logical Interface: 7F55C110
  Name: LOOPBACK                Protocol: LOOPBACK         State: Active
  Subnet Mask: 255.255.255.255  Addr: 127.0.0.1
  Mtu Size: 65535
Physical Interface: 7F3E2510
  Name: LOOPBACK6         Protocol: LOOPBACK6        State: Active
    NetNum: 0  QueSize: 0  Bytein: 0          Byteout: 0
     Index: 3
  Bsd Routing Parameters:
   MtuSize: 0                   Metric: 0
   SubnetMask: 0.0.0.0          DestAddr: 0.0.0.0
  SNMP Input Counters:
         Octets:                    0    Unicast:                    0
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Unkn Type:                    0
      Broadcast:                    0  Multicast:                    0
  SNMP Output Counters:
         Octets:                    0    Unicast:                    0
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Queue Len:                    0
      Broadcast:                    0  Multicast:                    0

  IPv4 Search Patricia tree     Address: 7F55CF10
    Search Ptree Reader Count: 0

  IPv6 Search Patricia tree     Address: 7F55C370
    Search Ptree Reader Count: 0

     Route: 7F3B4730
      Name: LOOPBACK6         Type: Host             State: Active
      Subnet Prefix: 128            Addr: ::1
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 65535         Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
```

```
      Metric5:  -1            Age: 2005/08/24 20:21:23


Logical Interface: 7F3B4A90
  Name: LOOPBACK6               Protocol: LOOPBACK6        State: Active
  Subnet Prefix: 128            Addr: ::1
  Mtu Size: 65535
================================================================================

Device Interface: 7F271410
  Device: VIPA16           Devtype: VIRTual          State: Active
  Address: **** ****

Physical Interface: 7F3E2D10
  Name: VIPA16            Protocol: VIRTUAL6         State: Active
    NetNum: 0  QueSize: 0  Bytein: 0          Byteout: 0
     Index: 5
  Bsd Routing Parameters:
   MtuSize: 0                   Metric: 0
   SubnetMask: 0.0.0.0          DestAddr: 0.0.0.0
  SNMP Input Counters:
         Octets:                    0    Unicast:                   0
     NonUnicast:                    0  Discarded:                   0
          Error:                    0  Unkn Type:                   0
      Broadcast:                    0  Multicast:                   0
  SNMP Output Counters:
         Octets:                    0    Unicast:                   0
     NonUnicast:                    0  Discarded:                   0
          Error:                    0  Queue Len:                   0
      Broadcast:                    0  Multicast:                   0

  IPv4 Search Patricia tree     Address: 7F55CF10
    Search Ptree Reader Count: 0
  IPv6 Search Patricia tree     Address: 7F55C370
    Search Ptree Reader Count: 0

     Route: 7F3D32B0
      Name: VIPA16            Type: Host             State: Active
      Subnet Prefix: 128            Addr: 50c9:c2d4::a:9:42:130:161
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 65535         Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:23



Logical Interface: 7F2791A8
  Name: VIPA16                  Protocol: VIRTUAL6         State: Active
  Subnet Prefix: 0              Addr: 50c9:c2d4::a:9:42:130:161
  Mtu Size: 65535
================================================================================
Device Interface: 7F26F410
  Device: IUTSAMEH         Devtype: MPCPTP           State: Setup
  Address: **** ****
SAP:
  UserID: 10010000     TransId: 00010140  ProviderId: 00010148
  Data@: 85B99378   ReqSignal@: 85B99378  RspSignal@: 85B99378
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0

Physical Interface: 7F3E3110
  Name: IUTSAMEH6          Protocol: MPCPTP6          State: Inactive
    NetNum: 0  QueSize: 0  Bytein: 0          Byteout: 0
     Index: 7
  Bsd Routing Parameters:
   MtuSize: 0                   Metric: 0
   SubnetMask: 0.0.0.0          DestAddr: 0.0.3.232
```

```
  SNMP Input Counters:
        Octets:                    0    Unicast:                   0
     NonUnicast:                   0  Discarded:                   0
          Error:                   0  Unkn Type:                   0
      Broadcast:                   0  Multicast:                   0
  SNMP Output Counters:
        Octets:                    0    Unicast:                   0
     NonUnicast:                   0  Discarded:                   0
          Error:                   0  Queue Len:                   0
      Broadcast:                   0  Multicast:                   0

  IPv4 Search Patricia tree     Address: 7F55CF10
    Search Ptree Reader Count: 0
  IPv6 Search Patricia tree     Address: 7F55C370
    Search Ptree Reader Count: 0

     Route: 7F2697F0
      Name: IUTSAMEH6         Type: Host             State: Inactive
      Subnet Prefix: 128            Addr: fe80::b47d:2e3f:c8c2:9117
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 65535         Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:23

     Route: 7F3D6448
      Name: IUTSAMEH6         Type: Host             State: Inactive
      Subnet Prefix: 128            Addr: fec0::42:105:75:161
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 65535         Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:23
Logical Interface: 7F269AD0
  Name: IUTSAMEH6               Protocol: MPCPTP6          State: Inactive
  Subnet Prefix: 128            Addr: fe80::b47d:2e3f:c8c2:9117
  Mtu Size: 65535
Logical Interface: 7F27A100
  Name: IUTSAMEH6               Protocol: MPCPTP6          State: Inactive
  Subnet Prefix: 0              Addr: fec0::42:105:75:161
  Mtu Size: 65535
================================================================================
Device Interface: 7F26E410
  Device: OSAQDIO2         Devtype: MPCIPA           State: Active
  Address: **** ****
SAP:
  UserID: 10020000     TransId: 00010141  ProviderId: 00010145
  Data@: 94EE174C   ReqSignal@: 85B99378  RspSignal@: 85B99378
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0
Connection 2:
  UserID: 00000000 ProviderId: 90020001
  Data@: 00000000  ReqSignal@: 00000000 RspSignal@: 00000000
  State: Reset        linknum: 00 flags 00

Physical Interface: 7F3E3510
  Name: OSAQDIO26         Protocol: IPAQENET6        State: Active
    NetNum: 0  QueSize: 0  Bytein: 1,008      Byteout: 1,304
     Index: 9
  Bsd Routing Parameters:
   MtuSize: 0                   Metric: 0
   SubnetMask: 0.0.0.0          DestAddr: 0.0.3.232
  SNMP Input Counters:
        Octets:                1,008    Unicast:                   0
     NonUnicast:                   0  Discarded:                   0
          Error:                   0  Unkn Type:                   0
      Broadcast:                   0  Multicast:                   6
  SNMP Output Counters:
```

```
        Octets:                1,304    Unicast:                    3
    NonUnicast:                    0  Discarded:                    0
         Error:                    0  Queue Len:                    0
     Broadcast:                    0  Multicast:                   10

  IPv4 Search Patricia tree     Address: 7F55CF10
    Search Ptree Reader Count: 0
  IPv6 Search Patricia tree     Address: 7F55C370
    Search Ptree Reader Count: 0

     Route: 7F24D270
      Name: OSAQDIO26         Type: Host             State: Active
      Subnet Prefix: 128            Addr: fe80::9:6b00:f71a:422
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 9000          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:23

     Route: 7F1C2090
      Name: OSAQDIO26         Type: Host             State: Active
      Subnet Prefix: 128            Addr: 50c9:c2d4::1:9:42:105:153
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 9000          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:27

     Route: 7F1D0BB0
      Name: OSAQDIO26         Type: Direct           State: Active
      Subnet Prefix: 64             Addr: 50c9:c2d4::1:0:0:0:0
      Protocol   : ICMP             Gate: ::0
      Mtu Size: 9000          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:25
     Route: 7F1D0830
      Name: OSAQDIO26         Type: Direct           State: Active
      Subnet Prefix: 64             Addr: 50c9:c2d4::1a:0:0:0:0
      Protocol   : ICMP             Gate: ::0
      Mtu Size: 9000          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:25



Logical Interface: 7F24D550
  Name: OSAQDIO26               Protocol: IPAQENET6        State: Active
  Subnet Prefix: 128            Addr: fe80::9:6b00:f71a:422
  Mtu Size: 9000
Logical Interface: 7F3B4190
  Name: OSAQDIO26               Protocol: IPAQENET6        State: Active
  Subnet Prefix: 0              Addr: 50c9:c2d4::1:9:42:105:153
  Mtu Size: 9000
================================================================================
Device Interface: 7F26D410
  Device: OSAQDIO4         Devtype: MPCIPA           State: Active
  Address: **** ****
SAP:
  UserID: 10030000     TransId: 00010142  ProviderId: 00010146
  Data@: 94EE174C   ReqSignal@: 85B99378  RspSignal@: 85B99378
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0
Connection 2:
  UserID: 00000000 ProviderId: 90030001
  Data@: 00000000  ReqSignal@: 00000000 RspSignal@: 00000000
   State: Reset        linknum: 00 flags 00
```

```
Physical Interface: 7F3E3910
  Name: OSAQDIO46         Protocol: IPAQENET6        State: Active
    NetNum: 0  QueSize: 0  Bytein: 3,120      Byteout: 1,444
     Index: 11
  Bsd Routing Parameters:
   MtuSize: 0                   Metric: 0
   SubnetMask: 0.0.0.0          DestAddr: 0.0.3.232
  SNMP Input Counters:
         Octets:                3,120    Unicast:                    2
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Unkn Type:                    0
      Broadcast:                    0  Multicast:                   19
  SNMP Output Counters:
         Octets:                1,444    Unicast:                    4
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Queue Len:                    0
      Broadcast:                    0  Multicast:                   10
  IPv4 Search Patricia tree     Address: 7F55CF10
    Search Ptree Reader Count: 0

  IPv6 Search Patricia tree     Address: 7F55C370
    Search Ptree Reader Count: 0

     Route: 7F2214E8
      Name: OSAQDIO46         Type: Host             State: Active
      Subnet Prefix: 128            Addr: fe80::9:6b01:f1a:684
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 1500          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:24

     Route: 7F1C7A10
      Name: OSAQDIO46         Type: Host             State: Active
      Subnet Prefix: 128            Addr: 50c9:c2d4::9:42:105:75
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 1500          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:29
                                                 Route: 7F3D3730
      Name: OSAQDIO46         Type: Direct           State: Active
      Subnet Prefix: 64             Addr: 50c9:c2d4::0
      Protocol   : Configuration    Gate: ::0
      Mtu Size: 1492          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:23

     Route: 7F1C7870
      Name: OSAQDIO46         Type: Direct           State: Active
      Subnet Prefix: 64             Addr: 50c9:c2d4::a:0:0:0:0
      Protocol   : ICMP             Gate: ::0
      Mtu Size: 1500          Ref Cnt: 0          Tos: 0
      Metric1:  0             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:29

     Route: 7F3D3590
      Name: OSAQDIO46         Type: Default          State: Active
      Subnet Prefix: 0              Addr: ::0
      Protocol   : Configuration    Gate: 50c9:c2d4::206:2aff:fe71:4400
      Mtu Size: 1492          Ref Cnt: 0          Tos: 0
      Metric1:  1             Meric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/24 20:21:23
```

```
 Logical Interface: 7F2217C8
   Name: OSAQDIO46               Protocol: IPAQENET6        State: Active
   Subnet Prefix: 128            Addr: fe80::9:6b01:f1a:684
   Mtu Size: 1500
 Logical Interface: 7F3B4050
  Name: OSAQDIO46               Protocol: IPAQENET6        State: Active
  Subnet Prefix: 0              Addr: 50c9:c2d4::9:42:105:75
  Mtu Size: 1500
================================================================================
Device Interface: 7F26C410
  Device: OSAQDIO7         Devtype: MPCIPA           State: Active
  Address: **** ****
SAP:
  UserID: 10040000     TransId: 00010143  ProviderId: 00010147
  Data@: 94EE174C   ReqSignal@: 85B99378  RspSignal@: 85B99378
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0
Connection 2:
  UserID: 00000000 ProviderId: 90040001
  Data@: 00000000  ReqSignal@: 00000000 RspSignal@: 00000000
  State: Reset        linknum: 00 flags 00
Physical Interface: 7F3E3D10
  Name: OSAQDIO76         Protocol: IPAQENET6        State: Active
    NetNum: 0  QueSize: 0  Bytein: 3,120      Byteout: 1,430
     Index: 13
  Bsd Routing Parameters:
   MtuSize: 0                   Metric: 0
   SubnetMask: 0.0.0.0          DestAddr: 0.0.3.232
  SNMP Input Counters:
         Octets:                3,120    Unicast:                    3
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Unkn Type:                    0
      Broadcast:                    0  Multicast:                   18
  SNMP Output Counters:
         Octets:                1,430    Unicast:                    3
     NonUnicast:                    0  Discarded:                    0
          Error:                    0  Queue Len:                    0
      Broadcast:                    0  Multicast:                   11

  IPv4 Search Patricia tree     Address: 7F55CF10
    Search Ptree Reader Count: 0
  IPv6 Search Patricia tree     Address: 7F55C370
    Search Ptree Reader Count: 0

      Route: 7F1C7C90
       Name: OSAQDIO76         Type: Host             State: Active
       Subnet Prefix: 128            Addr: 50c9:c2d4::9:42:105:85
       Protocol   : Configuration    Gate: ::0
       Mtu Size: 1500          Ref Cnt: 0          Tos: 0
       Metric1:  0             Meric2: -1
       Metric3:  -1            Metric4: -1
       Metric5:  -1            Age: 2005/08/24 20:21:28

      Route: 7F1C26F0
       Name: OSAQDIO76         Type: Direct           State: Active
       Subnet Prefix: 64             Addr: 50c9:c2d4::a:0:0:0:0
       Protocol   : ICMP             Gate: ::0
       Mtu Size: 1500          Ref Cnt: 0          Tos: 0
       Metric1:  0             Meric2: -1
       Metric3:  -1            Metric4: -1
       Metric5:  -1            Age: 2005/08/24 20:21:26

      Route: 7F21B190
       Name: OSAQDIO76         Type: Host             State: Active
       Subnet Prefix: 128            Addr: fe80::9:6b00:151a:594
```

```
       Protocol   : Configuration    Gate: ::0
       Mtu Size: 1500          Ref Cnt: 1          Tos: 0
       Metric1:  0             Meric2: -1
       Metric3:  -1            Metric4: -1
       Metric5:  -1            Age: 2005/08/24 20:21:24


 Logical Interface: 7F21B470
   Name: OSAQDIO76               Protocol: IPAQENET6        State: Active
   Subnet Prefix: 128            Addr: fe80::9:6b00:151a:594
   Mtu Size: 1500
 Logical Interface: 7F3D38D0
   Name: OSAQDIO76               Protocol: IPAQENET6        State: Active
   Subnet Prefix: 0              Addr: 50c9:c2d4::9:42:105:85
   Mtu Size: 1500
================================================================================
================================================================================

No IPv4 Lan Groups

IPv6 LAN Group Summary
LanGroup: 1     7F1C24B0
  IntfName           IntfStatus        NDOwner              VipaOwner
  --------           ----------        -------              ---------
  OSAQDIO76          Active            OSAQDIO76            Yes
  OSAQDIO46          Active            OSAQDIO46            No
LanGroup: 2     7F26F030
  IntfName           IntfStatus        NDOwner              VipaOwner
  --------           ----------        -------              ---------
  OSAQDIO26          Active            OSAQDIO26            Yes
Analysis of Tcp/Ip for TCPCS completed
________________________________________________________________________________
===== example output for a stack which is not IPv6 enabled ====================

Dataset: IPCS.X370812.TRAK0004.STEP2.DUMP
 Title:   TRAK0004 BEFORE


 The address of the TSAB is: 1524C000

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts          Status

 1524C040  1 TCPCS     V1R8    1500E000 1500E0C8 0030 9FFF767F 00000000  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V1R8  found

 ================================================================================

 Analysis of Tcp/Ip for TCPCS.  Index: 1

 TCPIP State

 TCPIP Status:
  Procedure:     TCPCS
  Version:       V1R8
  Status:        Active
  Asid:          0030
  Started:       2005/08/23 13:49:24
  Ended:         2005/08/23 15:01:12
  Active:        01:11:47.721301 hours

 Major Control Blocks
  TSEB:          1524C040       TSDB:          1500E000
```

```
 TSDX:          1500E0C8       TCA:           15B6B418
 ITCVT:         1500E398       ITSTOR:        1500E608
 DUAF:          1506B018       MRCB:          7F44F0F0
 MTCB:          15BB5B08       MUCB:          7F4250F0
 IPMAIN:        1505D398       Streams_root:  7F604D10
 TosMains:      15BB5948       MIB2:          15B6B078
 CdCb:          15D83510       User:          15EC2000
 Conf:          15BB53C8       Stks:          15D83450
================================================================================

TCPIP Subtasks
Task      Tcb       FirstRB   EotECB    StopEcb   CmpCode   RsnCode   RTWA
EZBTCPIP  008FF2A0  008FF218  808FD2C8            00000000  00000000  00000000
EPWPITSK  008E4E88  008DC110  00000000            00000000  00000000  00000000
EZBITTUB  008E4CF0  008EB340  00000000  808DC198  00000000  00000000  00000000
EZACDMSM  008E4B58  008FF4A8  00000000  808FF4A8  00000000  00000000  00000000
EZBIPSUB  008E49C0  008EB608  00000000  808DC088  00000000  00000000  00000000
EZBIEOER  008E4720  008E4698  008FF218  808E4698  940C4000  00000004  00000000
EZBTLMST  008E4500  008E4478  008FF218            00000000  00000000  00000000
EZACFMMN  008E41E8  008E4160  808FF218  808E4160  00000000  00000000  00000000
EZBTZMST  008EB470  008E43E0  008FF218            00000000  00000000  00000000
EZBTTSSL  008DCD90  008DCD08  808E43E0            00000000  00000000  00000000
EZBTMCTL  008DCA78  008DCF68  808E43E0            00000000  00000000  00000000
EZACFALG  008DC8E0  008DC858  808FF218            00000000  00000000  00000000
EZASASUB  008DEE88  008DCC70  808FF218  808DCC70  00000000  00000000  00000000
================================================================================
Storage Cache Information
  Total CSA Allocated:                7,703,656
  Tcp/ip CSA Limit:                      2,147M
  Total CSA Elements:                        47
  Cache Delay:                              300 seconds
  Scan Delay:                               120 seconds
  Total cache allocated:                 91,760
  Total cache elements:                       9
  Total freed elements:                       0
  Last cache scan time:     2005/08/23 18:59:41

CSM Status
  ECSA Storage:           OK
  Data Space Storage:     OK
  Fixed Storage:          OK
Alet: 01FF0014               Dspname: CSM64001
Alet: 00000000               Dspname: ........
================================================================================
Dispatchable Unit Status
  DUCB Initializations:                 64,057
  DUCB Expansions:                       3,684
  Percent DUCB expansions:                   5 %
  Last DUCB scan time:     2005/08/23 19:00:44

   1 DUAT control block(s) were found in the DUAF at 1506B018
 124 Dispatchable units were found.
  No DUs indicate abend.
================================================================================

CTrace Status:
   Member Name  : CTIEZBN0
   Buffer Size  : 4,194,304
   Options      : Init Opcmds Opmsgs Socket AFP XCF Access PFS
                  API Engine Queue RAW UDP TCP ICMP ARP ND CLAW
                  LCS Internet Message WorkUnit Config SNMP
                  IOCTL FireWall VtamData TelnVtam Telnet Vtam
   Asid List    : ()
   JobNameList  : ()
   PortList     : ()
   IpAddrList   : ()
   Xwriter      : Disconnected
   Dwriter      : Disconnected
```

```
   Trace Count  : 409,675
   Lost Count   : 2
   Lost Time    : 2005/08/23 17:49:25
   Wrap Count   : 35
   Wrap Time    : 2005/08/23 18:59:56
================================================================================

Device Interface: 7F607410
  Device: LOOPBACK         Devtype: LOOPBACK         State: Active
  Address: **** ****

Physical Interface: 7F604230
  Name: LOOPBACK          Protocol: LOOPBACK         State: Active
    NetNum: 0  QueSize: 0  Bytein: 43,385     Byteout: 43,385
     Index: 2
  Bsd Routing Parameters:
   SubnetMask: 0.0.0.0          DestAddr: 0.0.0.0
  SNMP Input Counters:
        Octets:               43,385    Unicast:                  677
     NonUnicast:                   0  Discarded:                    0
          Error:                   0  Unkn Type:                    0
      Broadcast:                   0  Multicast:                    0
  SNMP Output Counters:
        Octets:               43,385    Unicast:                  677
     NonUnicast:                   0  Discarded:                    0
          Error:                   0  Queue Len:                    0
      Broadcast:                   0  Multicast:                    0

  IPv4 Search Patricia tree     Address: 7F58EF70
    Search Ptree Reader Count: 0

     Route: 7F607270
      Name: LOOPBACK          Type: Host             State: Active
      Subnet Mask: 255.255.255.255  Addr: 127.0.0.1
      Protocol   : Configuration    Gate: 0.0.0.0
      Mtu Size: 65535         Ref Cnt: 5          Tos: 0
      Metric1:  0             Metric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/23 17:49:28



Logical Interface: 7F58E110
  Name: LOOPBACK                Protocol: LOOPBACK         State: Active
  Subnet Mask: 255.255.255.255  Addr: 127.0.0.1
  Mtu Size: 65535
================================================================================
Device Interface: 7F3DF410
  Device: IUTSAMEH         Devtype: MPCPTP           State: Setup
  Address: **** ****
SAP:
  UserID: 10010000     TransId: 00010130  ProviderId: 00010136
  Data@: 83BED5E0   ReqSignal@: 83BED5E0  RspSignal@: 83BED5E0
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0

Physical Interface: 7F414510
  Name: LSAMEH            Protocol: MPCPTP           State: Inactive
    NetNum: 0  QueSize: 0  Bytein: 0          Byteout: 0
     Index: 4
  Bsd Routing Parameters:
   MtuSize: 576                 Metric: 0
   SubnetMask: 255.0.0.0        DestAddr: 0.0.0.0
  SNMP Input Counters:
        Octets:                   0    Unicast:                    0
     NonUnicast:                  0  Discarded:                    0
          Error:                  0  Unkn Type:                    0
      Broadcast:                  0  Multicast:                    0
```

```
  SNMP Output Counters:
        Octets:                    0    Unicast:                    0
     NonUnicast:                   0  Discarded:                    0
          Error:                   0  Queue Len:                    0
      Broadcast:                   0  Multicast:                    0
  IPv4 Search Patricia tree     Address: 7F58EF70
    Search Ptree Reader Count: 0

     Route: 7F4082E8
      Name: LSAMEH            Type: Host            State: Inactive
      Subnet Mask: 255.255.255.255  Addr: 10.1.0.161
      Protocol   : Configuration    Gate: 0.0.0.0
      Mtu Size: 65535         Ref Cnt: 0          Tos: 0
      Metric1:  0             Metric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/23 17:49:32



Logical Interface: 7F3DF090
  Name: LSAMEH                  Protocol: MPCPTP           State: Inactive
  Subnet Mask: 255.255.255.255  Addr: 10.1.0.161
  Mtu Size: 65535
================================================================================
Device Interface: 7F3DD410
  Device: OSAQDIO2         Devtype: MPCIPA           State: Active
  Address: **** ****
SAP:
  UserID: 30040000     TransId: 00010145  ProviderId: 00010147
  Data@: 9500474C   ReqSignal@: 83BED5E0  RspSignal@: 83BED5E0
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0
Connection 2:
  UserID: 00000000 ProviderId: 90040001
  Data@: 00000000  ReqSignal@: 00000000 RspSignal@: 00000000
  State: Reset        linknum: 00 flags 00

Physical Interface: 7F414910
  Name: LOSAQDIO2         Protocol: IPAQENET         State: Active
    NetNum: 0  QueSize: 0  Bytein: 297,780    Byteout: 32,736
     Index: 6
  Bsd Routing Parameters:
   MtuSize: 576                 Metric: 1
   SubnetMask: 255.255.255.128  DestAddr: 0.0.0.0
  SNMP Input Counters:
        Octets:              297,780    Unicast:                    6
     NonUnicast:                   0  Discarded:                    0
          Error:                   0  Unkn Type:                    0
      Broadcast:                  66  Multicast:                2,027
  SNMP Output Counters:
        Octets:               32,736    Unicast:                   10
     NonUnicast:                   0  Discarded:                    0
          Error:                   0  Queue Len:                    0
      Broadcast:                   0  Multicast:                  218

  IPv4 Search Patricia tree     Address: 7F58EF70
    Search Ptree Reader Count: 0

     Route: 7F66ECD0
      Name: LOSAQDIO2         Type: Host            State: Active
      Subnet Mask: 255.255.255.255  Addr: 9.42.130.173
      Protocol   : OSPF             Gate: 9.42.105.139
      Mtu Size: 576           Ref Cnt: 0          Tos: 0
      Metric1:  2             Metric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/23 18:31:15
```

```
Route: 7F66E990
 Name: LOSAQDIO2         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.130.134
 Protocol   : OSPF             Gate: 9.42.105.158
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  2             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:31:15

Route: 7F66EE70
 Name: LOSAQDIO2         Type: Subnetwork       State: Active
 Subnet Mask: 255.255.255.252  Addr: 9.42.130.172
 Protocol   : OSPF             Gate: 9.42.105.139
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  2             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:31:15

Route: 7F66E7F0
 Name: LOSAQDIO2         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.130.85
 Protocol   : OSPF             Gate: 9.42.105.136
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  2             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:31:15

Route: 7F66E650
 Name: LOSAQDIO2         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.130.46
 Protocol   : OSPF             Gate: 9.42.105.149
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  2             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:31:15

Route: 7F326AF0
 Name: LOSAQDIO2         Type: Subnetwork       State: Active
 Subnet Mask: 255.255.255.252  Addr: 9.42.130.44
 Protocol   : OSPF             Gate: 9.42.105.149
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  2             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:31:15

Route: 7F66EB30
 Name: LOSAQDIO2         Type: Subnetwork       State: Active
 Subnet Mask: 255.255.255.252  Addr: 9.42.130.132
 Protocol   : OSPF             Gate: 9.42.105.158
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  2             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:31:15

Route: 7F585310
 Name: LOSAQDIO2         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.105.153
 Protocol   : Configuration    Gate: 0.0.0.0
 Mtu Size: 576           Ref Cnt: 1          Tos: 0
 Metric1:  0             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 17:49:32

Route: 7F320D10
 Name: LOSAQDIO2         Type: Direct           State: Active
 Subnet Mask: 255.255.255.128  Addr: 9.42.105.128
 Protocol   : OSPF             Gate: 0.0.0.0
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
```

```
    Metric1:  1             Metric2: -1
    Metric3:  -1            Metric4: -1
    Metric5:  -1            Age: 2005/08/23 18:31:04

   Route: 7F326950
    Name: LOSAQDIO2         Type: Subnetwork       State: Active
    Subnet Mask: 255.255.255.128  Addr: 9.42.103.128
    Protocol   : OSPF             Gate: 9.42.105.129
    Mtu Size: 576           Ref Cnt: 0          Tos: 0
    Metric1:  2             Metric2: -1
    Metric3:  -1            Metric4: -1
    Metric5:  -1            Age: 2005/08/23 18:31:15


  Address Translate Entry: 7F31FB90
    addr: 9.42.105.153     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F31FC50
    addr: 9.42.105.184     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F31FD10
    addr: 9.42.105.143     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F31FDD0
    addr: 9.42.105.141     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F31FE90
    addr: 9.42.105.138     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F31FF50
    addr: 9.42.105.136     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F66E050
    addr: 9.42.105.130     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F66E110
    addr: 9.42.105.139     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F66E1D0
    addr: 9.42.105.133     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F66E290
    addr: 9.42.105.155     flags:  C0  ttldlt:  0
    retries: 0
  Address Translate Entry: 7F66E350
    addr: 9.42.105.179     flags:  C0  ttldlt:  0
    retries: 0
   Address Translate Entry: 7F66E410
     addr: 9.42.105.142     flags:  C0  ttldlt:  0
     retries: 0
   Address Translate Entry: 7F66E4D0
     addr: 9.42.105.144     flags:  C0  ttldlt:  0
     retries: 0
   Address Translate Entry: 7F322070
     addr: 9.42.105.195     flags:  C0  ttldlt:  0
     retries: 0
   Address Translate Entry: 7F322C70
     addr: 9.42.105.152     flags:  C0  ttldlt:  0
     retries: 0
   Address Translate Entry: 7F33AB50
     addr: 9.42.105.129     flags:  C0  ttldlt:  0
     retries: 0

Logical Interface: 7F3E71A8
  Name: LOSAQDIO2               Protocol: IPAQENET         State: Active
  Subnet Mask: 255.255.255.255  Addr: 9.42.105.153
  Mtu Size: 8992
```

```
================================================================================
Device Interface: 7F3DC410
  Device: OSAQDIO4         Devtype: MPCIPA           State: Active
  Address: **** ****
SAP:
  UserID: 10020000     TransId: 00010131  ProviderId: 00010134
  Data@: 9500474C   ReqSignal@: 83BED5E0  RspSignal@: 83BED5E0
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0
Connection 2:
  UserID: 00000000 ProviderId: 90020001
  Data@: 00000000  ReqSignal@: 00000000 RspSignal@: 00000000
  State: Reset        linknum: 00 flags 00

Physical Interface: 7F414D10
  Name: LOSAQDIO4         Protocol: IPAQENET         State: Active
    NetNum: 0  QueSize: 0  Bytein: 2,990      Byteout: 924
     Index: 8
  Bsd Routing Parameters:
   MtuSize: 576                Metric: 1
  SubnetMask: 255.255.255.128  DestAddr: 0.0.0.0
 SNMP Input Counters:
        Octets:                2,990    Unicast:                    1
    NonUnicast:                    0  Discarded:                    0
         Error:                    0  Unkn Type:                    0
     Broadcast:                   11  Multicast:                    0
 SNMP Output Counters:
        Octets:                  924    Unicast:                    3
    NonUnicast:                    0  Discarded:                    0
         Error:                    0  Queue Len:                    0
     Broadcast:                    0  Multicast:                    0

 IPv4 Search Patricia tree     Address: 7F58EF70
   Search Ptree Reader Count: 0

    Route: 7F334E70
     Name: LOSAQDIO4         Type: Host             State: Active
     Subnet Mask: 255.255.255.255  Addr: 9.42.130.117
     Protocol   : OSPF             Gate: 9.42.105.110
     Mtu Size: 576           Ref Cnt: 0          Tos: 0
     Metric1:  2             Metric2: -1
     Metric3:  -1            Metric4: -1
     Metric5:  -1            Age: 2005/08/23 18:16:51

    Route: 7F3D7170
     Name: LOSAQDIO4         Type: Host             State: Active
     Subnet Mask: 255.255.255.255  Addr: 9.42.130.47
     Protocol   : OSPF             Gate: 9.42.105.65
     Mtu Size: 576           Ref Cnt: 0          Tos: 0
     Metric1:  3             Metric2: -1
     Metric3:  -1            Metric4: -1
     Metric5:  -1            Age: 2005/08/23 18:16:51

    Route: 7F338290
     Name: LOSAQDIO4         Type: Host             State: Active
     Subnet Mask: 255.255.255.255  Addr: 9.42.130.11
     Protocol   : OSPF             Gate: 9.42.105.65
     Mtu Size: 576           Ref Cnt: 0          Tos: 0
     Metric1:  3             Metric2: -1
     Metric3:  -1            Metric4: -1
     Metric5:  -1            Age: 2005/08/23 18:16:51

    Route: 7F338DD0
     Name: LOSAQDIO4         Type: Subnetwork       State: Active
     Subnet Mask: 255.255.255.252  Addr: 9.42.130.8
     Protocol   : OSPF             Gate: 9.42.105.65
     Mtu Size: 576           Ref Cnt: 0          Tos: 0
```

```
 Metric1:  3             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:51

Route: 7F3D75F0
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.130.48
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  3             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:51

Route: 7F3992F0
 Name: LOSAQDIO4         Type: Subnetwork       State: Active
 Subnet Mask: 255.255.255.252  Addr: 9.42.130.48
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  3             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:51

Route: 7F338850
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.130.12
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  3             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:51

Route: 7F3DA3F0
 Name: LOSAQDIO4         Type: Subnetwork       State: Active
 Subnet Mask: 255.255.255.252  Addr: 9.42.130.12
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  3             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:51

Route: 7F585030
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.105.75
 Protocol   : Configuration    Gate: 0.0.0.0
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  0             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 17:49:32

Route: 7F343070
 Name: LOSAQDIO4         Type: Direct           State: Active
 Subnet Mask: 255.255.255.128  Addr: 9.42.105.0
 Protocol   : OSPF             Gate: 0.0.0.0
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:51

Route: 7F329AD0
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.188
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F329450
```

```
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.189
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

 Route: 7F3227B0
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.186
 Protocol   : OSPF             Gate: 9.42.105.63
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32A150
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.185
 Protocol   : OSPF             Gate: 9.42.105.63
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F322950
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.187
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32A930
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.176
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32A430
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.180
 Protocol   : OSPF             Gate: 9.42.105.63
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32B930
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.170
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32C290
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.168
 Protocol   : OSPF             Gate: 9.42.105.126
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
```

```
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32BE70
 Name: LOSAQDIO4         Type: Host            State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.169
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32B3F0
 Name: LOSAQDIO4         Type: Host            State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.171
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32CE70
 Name: LOSAQDIO4         Type: Host            State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.162
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F38D108
 Name: LOSAQDIO4         Type: Host            State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.161
 Protocol   : OSPF             Gate: 9.42.105.45
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  2             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:51

Route: 7F32D470
 Name: LOSAQDIO4         Type: Host            State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.160
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32C9F0
 Name: LOSAQDIO4         Type: Host            State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.163
  Protocol   : OSPF             Gate: 9.42.105.65
  Mtu Size: 576           Ref Cnt: 0          Tos: 0
  Metric1:  1             Metric2: -1
  Metric3:  -1            Metric4: -1
  Metric5:  -1            Age: 2005/08/23 18:16:52

 Route: 7F32C570
  Name: LOSAQDIO4         Type: Host            State: Active
  Subnet Mask: 255.255.255.255  Addr: 9.42.104.166
  Protocol   : OSPF             Gate: 9.42.105.121
  Mtu Size: 576           Ref Cnt: 0          Tos: 0
  Metric1:  1             Metric2: -1
  Metric3:  -1            Metric4: -1
  Metric5:  -1            Age: 2005/08/23 18:16:52

 Route: 7F32AE70
  Name: LOSAQDIO4         Type: Host            State: Active
```

```
  Subnet Mask: 255.255.255.255  Addr: 9.42.104.172
  Protocol   : OSPF             Gate: 9.42.105.65
  Mtu Size: 576           Ref Cnt: 0          Tos: 0
  Metric1:  1             Metric2: -1
  Metric3:  -1            Metric4: -1
  Metric5:  -1            Age: 2005/08/23 18:16:52

 Route: 7F399810
  Name: LOSAQDIO4         Type: Host             State: Active
  Subnet Mask: 255.255.255.255  Addr: 9.42.104.159
  Protocol   : OSPF             Gate: 9.42.105.45
  Mtu Size: 576           Ref Cnt: 0          Tos: 0
  Metric1:  51            Metric2: -1
  Metric3:  -1            Metric4: -1
  Metric5:  -1            Age: 2005/08/23 18:16:51

 Route: 7F32E570
  Name: LOSAQDIO4         Type: Host             State: Active
  Subnet Mask: 255.255.255.255  Addr: 9.42.104.156
  Protocol   : OSPF             Gate: 9.42.105.65
  Mtu Size: 576           Ref Cnt: 0          Tos: 0
  Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32DE70
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.157
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32EE70
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.154
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32F4F0
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.152
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32E9F0
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.155
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
 Metric5:  -1            Age: 2005/08/23 18:16:52

Route: 7F32FE70
 Name: LOSAQDIO4         Type: Host             State: Active
 Subnet Mask: 255.255.255.255  Addr: 9.42.104.150
 Protocol   : OSPF             Gate: 9.42.105.65
 Mtu Size: 576           Ref Cnt: 0          Tos: 0
 Metric1:  1             Metric2: -1
 Metric3:  -1            Metric4: -1
```

```
      Metric5:  -1            Age: 2005/08/23 18:16:52

     Route: 7F3224D0
      Name: LOSAQDIO4         Type: Host            State: Active
      Subnet Mask: 255.255.255.255  Addr: 9.42.104.149
      Protocol   : OSPF             Gate: 9.42.105.126
      Mtu Size: 576           Ref Cnt: 0          Tos: 0
      Metric1:  1             Metric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/23 18:16:52

     Route: 7F330190
      Name: LOSAQDIO4         Type: Host            State: Active
      Subnet Mask: 255.255.255.255  Addr: 9.42.104.146
      Protocol   : OSPF             Gate: 9.42.105.65
      Mtu Size: 576           Ref Cnt: 0          Tos: 0
      Metric1:  1             Metric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/23 18:16:52

     Route: 7F32F9F0
      Name: LOSAQDIO4         Type: Host            State: Active
      Subnet Mask: 255.255.255.255  Addr: 9.42.104.151
      Protocol   : OSPF             Gate: 9.42.105.65
      Mtu Size: 576           Ref Cnt: 0          Tos: 0
      Metric1:  1             Metric2: -1
      Metric3:  -1            Metric4: -1
      Metric5:  -1            Age: 2005/08/23 18:16:52

     Route: 7F330C90
      Name: LOSAQDIO4         Type: Host            State: Active
      Subnet Mask: 255.255.255.255  Addr: 9.42.104.142
     addr: 9.42.105.75      flags:  C0  ttldlt:  0
     retries: 0

Logical Interface: 7F3E7068
  Name: LOSAQDIO4               Protocol: IPAQENET         State: Active
  Subnet Mask: 255.255.255.255  Addr: 9.42.105.75
  Mtu Size: 1492
================================================================================
Device Interface: 7F3DB410
  Device: OSAQDIO7         Devtype: MPCIPA           State: Active
  Address: **** ****
SAP:
  UserID: 10030000     TransId: 00010132  ProviderId: 00010135
  Data@: 9500474C   ReqSignal@: 83BED5E0  RspSignal@: 83BED5E0
  State: Unknown        Retry:    0  Restart:    0  Xstatus: 0
Connection 2:
  UserID: 00000000 ProviderId: 90030001
  Data@: 00000000  ReqSignal@: 00000000 RspSignal@: 00000000
  State: Reset        linknum: 00 flags 00

Physical Interface: 7F415110
  Name: LOSAQDIO7         Protocol: IPAQENET         State: Active
    NetNum: 0  QueSize: 0  Bytein: 583,077    Byteout: 76,876
     Index: 10
  Bsd Routing Parameters:
   MtuSize: 576                 Metric: 1
   SubnetMask: 255.255.255.128  DestAddr: 0.0.0.0
  SNMP Input Counters:
        Octets:              583,077    Unicast:                  36
    NonUnicast:                    0  Discarded:                   0
         Error:                    0  Unkn Type:                   0
     Broadcast:                   10  Multicast:               4,121
  SNMP Output Counters:
        Octets:               76,876    Unicast:                  38
    NonUnicast:                    0  Discarded:                   0
         Error:                    0  Queue Len:                   0
```

```
      Broadcast:                    0  Multicast:                  482

  IPv4 Search Patricia tree     Address: 7F58EF70
      Metric5:  -1            Age: 2005/08/23 18:16:52


   Address Translate Entry: 7F337030
     addr: 9.42.105.85      flags:  C0  ttldlt:  0
     retries: 0

Logical Interface: 7F3E8100
  Name: LOSAQDIO7               Protocol: IPAQENET         State: Active
  Subnet Mask: 255.255.255.255  Addr: 9.42.105.85
  Mtu Size: 1492
================================================================================
================================================================================


IPv4 LAN Group Summary
LanGroup: 1     7F3DABD0
  LnkName            LnkStatus         ArpOwner             VipaOwner
  -------            ---------         --------             ---------
  LOSAQDIO4          Active            LOSAQDIO4            Yes
  LOSAQDIO7          Active            LOSAQDIO7            No
LanGroup: 2     7F320B90
  LnkName            LnkStatus         ArpOwner             VipaOwner
  -------            ---------         --------             ---------
  LOSAQDIO2          Active            LOSAQDIO2            Yes

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS STORAGE

Use this subcommand to display the TCP/IP storage summary referenced in common cached storage.

Under the heading Storage Summary, a "c" in column "c" indicates the address is on the cache queue. A "p" in column "p" indicates that the control block is part of a pool.

Cache storage has 12 bytes from offset four overlaid with a chain pointer and time stamp. This might show incorrect data for cached control blocks.

Tip: The TCPIPCS STORAGE command only reports storage found in caches in common storage. Use the TCPIPCS MAP command to report both common and TCP/IP private storage usage.

Syntax

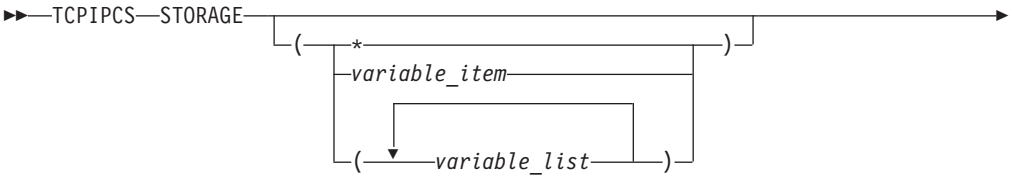

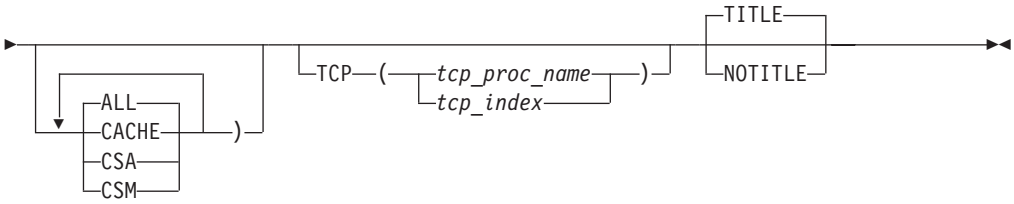

Parameters

ALL
Display information about all allocated storage.

CACHE
Display only information about cached storage.

CSA
Display only information about in-use CSA storage.

CSM
Display only information about in-use CSM storage.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1-32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

storage_address
Formats the storage headers for the TCPIP storage element at address storage_address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a-f or A-F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

When a hexadecimal address is specified as *variable_item* (or more than one address is specified as *variable_list*) the timestamp, module name, and an indication of whether the storage is allocated or freed, will be formatted.

Rule: If you specify multiple keywords from the set {ALL,CACHE,CSA,CSM}, all of the keywords that you specify are used.

When a BLS18100I message indicating an access failure appears in the report, any counts or analysis dependent on this information cannot be included in the TCPIPCS STORAGE output. Also, an access failure can occur as a result of insufficient user region size. If a BLS18100I message is received for data that is included in the dump, increase the user region size and attempt the TCPIPCS STORAGE subcommand again.

Sample output of the TCPIPCS STORAGE subcommand

The following is sample output of the TCPIPCS STORAGE subcommand:

```
TCPIPCS STORAGE
 Dataset: IPCS.A594094.DUMPM
 Title:   TCPSVT   V2R10: Job(TCPSVT  ) EZBITSTO(HTCP50A 99.281)+
          00077A S4C5/74BE2500 SRB P=0051,S=0051,H=0051
```

```
...
TCPIP Storage Analysis

Storage Statistics
cache_delay                    0 seconds before cache is freed
com_totstor          177,578,656 total storage for CSA elements
com_totelem               21,469 total number of CSA elements
scan_delay                   120 seconds between full scans
stor_cache                48,416 storage in cache after scan
num_cache                     11 elements in cache after scan
num_freed                      2 elements freed during last scan
scan_time    1999/10/24 04:06:12 time of last scan
dsa_init              10,375,262 # of DUCB initializations
dsa_exp                2,180,028 # of DUCB expansions
The control block at 008AC010 (Prev: 00000000) has already been added
...
The control block at 12A26410 (Prev: 137CB0A0) has already been added

 21,907 storage elements found
177,228K bytes of storage allocated

Cached Storage
Addr       Size Key Sp Cblk Time Stamp         Index

Common non-fetch protected storage
12E6DCB0    304  6 241 CFGM B30A8EDF19BD18C3  10
12774310   3056  6 241 CFGM B30A8E3DDBBB1943  10 Index was 29
The control block at 0E289010 (prev: 12B57650) was not available
Unable to locate storage at 0E289010
Cache pointers are in a loop at 12774310 for index 29
The control block at 0E289010 (prev: 12B57730) was not available
Unable to locate storage at 0E289010
      2 control blocks found for Common non-fetch protected storage
   3376 bytes allocated in Common non-fetch
4366931 total allocations

  Addr    Size Key Sp  Time Stamp (GMT)  Alloc Common PoolAddr  Module          Type
7F10A110  1264  6 249 2007/08/22 12:56:03  Y     N    7F10A010 959D13D6 private, non-fetch protected
7F115010  1264  6 249 2007/08/22 12:56:41  Y     N    7F10A010 959D13D6 private, non-fetch protected
7F112D10  1264  6 249 2007/08/22 13:52:06  Y     N    7F10A010 959D13D6 private, non-fetch protected
7F112310  1264  6 249 2007/08/22 13:52:11  Y     N    7F10A010 959D13D6 private, non-fetch protected
7F10CE10  1264  6 249 2007/08/22 12:58:15  Y     N    7F10A010 959D13D6 private, non-fetch protected
7F10C910  1264  6 249 2007/08/22 12:58:15  Y     N    7F10A010 959D13D6 private, non-fetch protected

...
Storage Summary Statistics
                                        All                Cache
Type                           Count      Size     Count      Size
Common Non-fetch protected     21460   177489K         2      3392
Common Fetch protected           369     68488       141     36936
Common persistent                  3       192         3       192
Common SCB pool                   80     21128        32      8448
Private Non-fetch protected      492    395848       156     65192
Total                          22571   178149K       334    114160

22599 blocks of storage for 1807728 bytes were obtained to create this report

Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS STREAM

Use this subcommand to display the stream control blocks.

Syntax

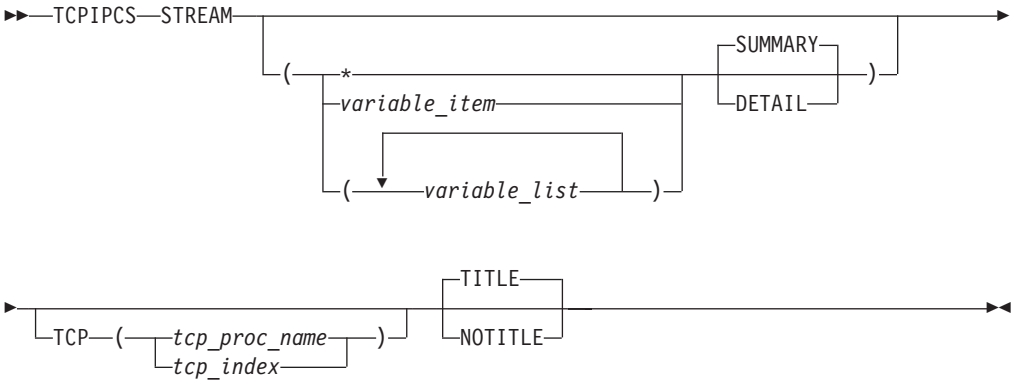

Parameters

If no parameters are specified, all stream control blocks are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

CB_address
An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string. Displays only the Stream control block associated with one of the following:

SKCB Stream context control block address.

SKQI Stream queue initialization control block address.

SKQP Stream queue pair control block address.

SKQU Stream Queue control block address.

SKSC Stream access control control block address.

SKSH Stream header control block address.

connection_id
Displays the Stream control block with this connection ID. A connection ID is specified as 1-8 hexadecimal digits.

In addition to the variable parameters described above, you can specify the following keyword parameters:

SUMMARY
Formats the Stream control blocks in one cross-reference table. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of the Stream control blocks.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS STREAM subcommand

The following is a sample output of the TCPIPCS STREAM subcommand:

```
TCPIPCS STREAM
 Dataset: IPCS.A594094.DUMPM
 Title:   TCPSVT   V2R10: Job(TCPSVT  ) EZBITSTO(HTCP50A 99.281)+
          00077A S4C5/74BE2500 SRB P=0051,S=0051,H=0051

 The address of the TSAB is: 12E89BB8

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 12E89BF8  1 TCPSVT    V2R10   12B57000 12B570C8 0051 9FFFFF7F  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPSVT.  Index: 1

 TCPIP Stream Analysis

 SKRT at 7F78BD88

 Sksc@    Sksh@    CID      Driver Api@     Skcb@    Ascb@    Tcb@
 7F77E6C8 7F77E7C8 00000007 IP/NAM 00000000 00000000 00000000 00000000
 7F70F088 7F61A088 00000006 RAW    00000000 00000000 00000000 00000000
 7F70F148 7F61A608 00000005 IP/NAM 00000000 00000000 00000000 00000000
 7F70F8C8 7F70F348 00000004 UDP    00000000 00000000 00000000 00000000
 7F70F988 7F70FA48 00000003 IP/NAM 00000000 00000000 00000000 00000000
 7F78B008 7F7580E8 00000002 TCP    00000000 00000000 00000000 00000000
 7F78BCC8 7F78B748 00000001 IP/NAM 00000000 00000000 00000000 00000000

 7 Stream(s) found
 7 Stream(s) formatted

 Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS TCB

Use this subcommand to display the Master Transmission Control Block (MTCB) and any Transmission protocol Control Blocks (TCBs) that are defined in the TCP hash table.

Syntax

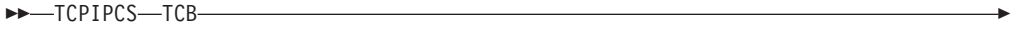

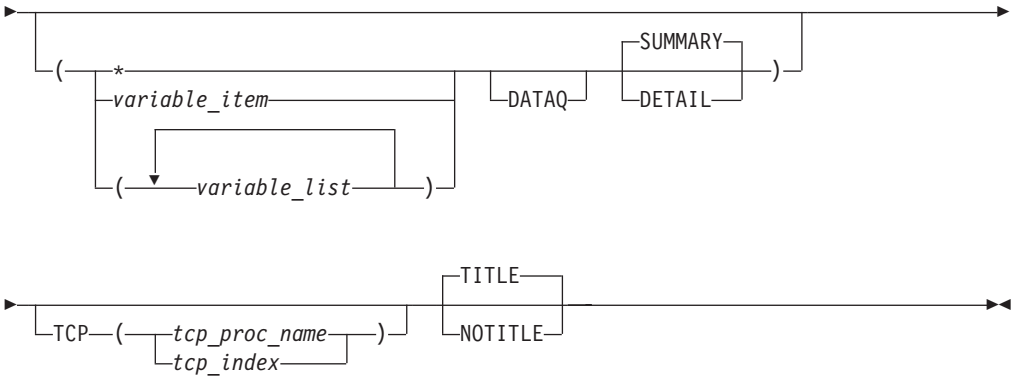

Parameters

If no parameters are specified, all TCP control blocks are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

jobname
Displays only the TCBs with this job name. The job name can be a TCP/IP application name or a stack name. A job name is 1–8 alphanumeric characters.

TCB_address
Displays only the TCB with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

connection_id
Displays the TCB with this connection ID. A connection ID is specified as 1-8 hexadecimal digits.

In addition to the variable parameters described above, you can specify the following keyword parameters:

DATAQ
Formats TCBs which have data queued on the SEND or RECEIVE queue.

SUMMARY
Formats the MTCB contents and lists all the TCBs in one cross-reference table. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of the TCBs.

TCP, TITLE, NOTITLE
See “Parameters” on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS TCB subcommand

The following is sample output of the TCPIPCS TCB subcommand:

```
TCPIPCS TCB
Dataset: IPCS.MV21372.DUMPA
Title:   SLIP DUMP ID=TC


The address of the TSAB is: 131B8120

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

131B8160  1 TCPSVT    V2R10   13C9F000 13C9F0C8 07D3 94FF755F  Active

   1 defined TCP/IP(s) were found
   1 active  TCP/IP(s) were found

   1 TCP/IP(s) for CS     V2R10 found

================================================================================

Analysis of Tcp/Ip for TCPSVT.  Index: 1


TCP/IP Analysis
TCPIP Main TCP Control Block (MTCB)
 MTCB: 13C9E890
 +0000  M_MAIN_EYE......... TCP MAIN
 +0008  M_TCP_LWRITE_Q..... 7F782868
 +000C  M_TCP_LREAD_Q...... 7F782828
 +0014  M_TCP_DRIVER_STATE. 01
 +0018  MTCPMTX............ 00000000  00000000  00000000  D7D60601
 +0028  MTCPAQMX........... 00000000  00000000  00000000  D7D60604
 +0038  MTCB_LIST_LOCK..... 00000000  00000000  00000000  D7D60604
 +0048  M_PORT_CEILING..... 00000FFF
 +004C  M_TPI_SEQ#......... 00000008
 +0050  M_PORT_ARRAY....... 7F711FC8
 +0054  M_LAST_PORT_NUM.... 0000040C
 ...

 TCB      ResrcID  ResrcNm  TcpState    TpiState Local IPAddr/Port    Remote IPAddr/Port    LuName   ApplName UserID
 7F603108 00000002 TCPSVT   Closed      WLOUNBND 0.0.0.0..0           0.0.0.0..0
 7F605D08 00000017 FTPUNIX1 Listening   WLOIDLE  0.0.0.0..21          0.0.0.0..0
 7F605108 00000013 TCPSVT   Listening   WLOIDLE  0.0.0.0..625         0.0.0.0..0
 7F603508 0000000A TCPSVT   Listening   WLOIDLE  0.0.0.0..1025        0.0.0.0..0
 7F604508 000000EA TCPSVT   Established WLOXFER  197.66.103.1..23     197.11.108.1..1032
 ...
 7F607108 0000003E TCPSVT   Established WLOXFER  127.0.0.1..1029      127.0.0.1..1028
 7F60A508 000000E8 TCPSVT   Listening   WLOIDLE  0.0.0.0..623         0.0.0.0..0
 25 TCB(s) FOUND
 25 TCB(s) FORMATTED

 Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS TELNET

Use this subcommand to display either the address, or address and contents, of Telnet control blocks. These include the following:

- TCMA
- TCFG
- TPDB
- Optionally, the TKCB and CVB for a selected session
- A partial TCFG that is being built is also displayed (if found)

Syntax

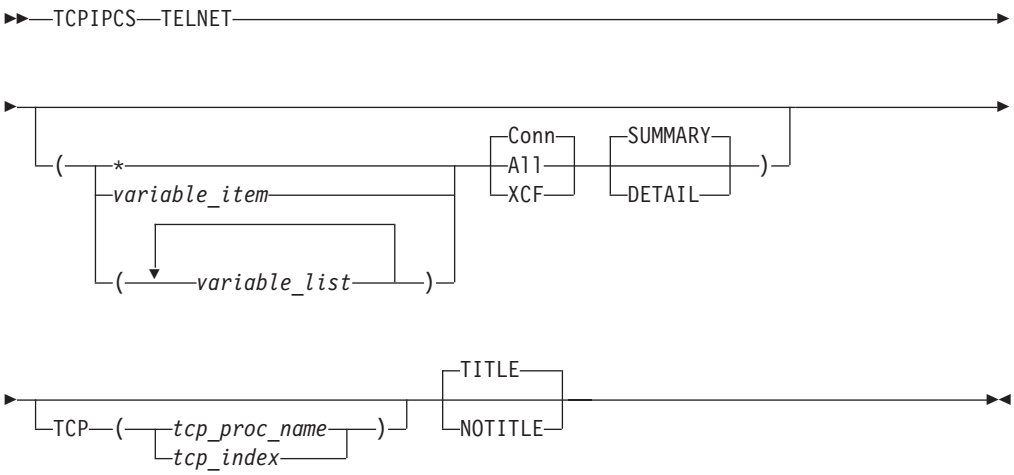

Parameters

If no parameters are specified, all TCP control blocks are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 – 32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

LUname
Displays only the session control blocks for the 8-character logical unit name. If the name is less than eight characters, it is padded on the right with blanks.

CVB_address
Displays only the CVB with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a-f or A-F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

token Displays only the session control blocks for the token. The token is a 16-digit hexadecimal value. If the token is less than 16 digits, it is padded on the left with hex zeros.

In addition to the variable parameters described above, you can specify the following keyword parameters:

All
Display Telnet connection and XCF information.

Conn
Display only Telnet connection information.

DETAIL
Displays the contents of the control blocks.

SUMMARY
Displays the address of the control blocks. SUMMARY is the default.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Xcf
Display only XCF information.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS TELNET subcommand

The following is sample output of the TCPIPCS TELNET subcommand:

```
TCPIPCS TELNET
Dataset: IPCS.MV21381.DUMPA
Title:   SLIP DUMP ID=TC


The address of the TSAB is: 13391BC0

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

13391C00  1 TCPSVT    V2R10   1323B000 1323B0C8 07DE 04041405  Active
13391C80  2 TCPSVT2   V2R10   00000000 00000000 07E8 00000000  Down Stopping
13391D00  3 TCPSVT1   V2R10   12FC3000 12FC30C8 0080 94FF755F  Active
13391D80  4 TCPSVT3   V2R10   00000000 00000000 0059 00000000  Down Stopping

   4 defined TCP/IP(s) were found
   2 active  TCP/IP(s) were found

   4 TCP/IP(s) for CS     V2R10 found

================================================================================

Analysis of Tcp/Ip for TCPSVT.  Index: 1

TCPIP Telnet Analysis

TMCA at 7F5B1188

Tpdb@     Port Tcfg@    Prof Tkcb@    Token             Cvb@     LUname
7F59D8A0   623 7F5A6068 CURR 00000000 00000000 00000000 00000000
7F59D4E0   625 7F59D620 CURR 00000000 00000000 00000000 00000000

Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS TIMER

Use this subcommand to display the timer control blocks.

Syntax

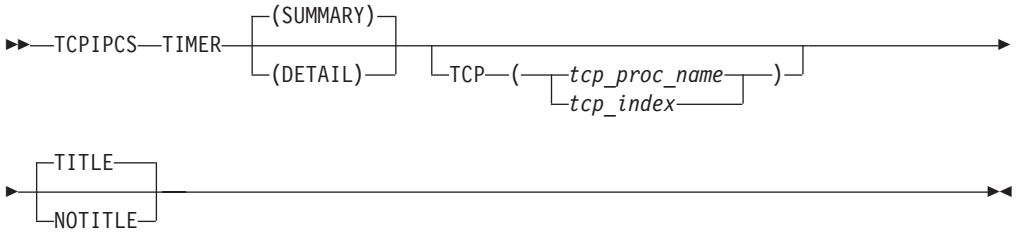

Parameters

SUMMARY
Displays the contents of the timer control blocks. The timer queue elements (TQEs) and timer IDs (TIDs) are presented in tabular form. SUMMARY is the default.

DETAIL
The timer control blocks are displayed as in the SUMMARY form of the command. In addition, each TQE and each TID is fully displayed.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS TIMER subcommand

The following is sample output of the TCPIPCS TIMER subcommand:

```
TCPIPCS TIMER
 Dataset: IPCS.A594094.DUMPF
 Title:   CHECK NOT ADDR


 The address of the TSAB is: 08CE28C0

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 08CE2900  1 TCPCS     V2R10   086D8000 086D80C8 01F8 10000100  Active

    1 defined TCP/IP(s) were found
    1 active  TCP/IP(s) were found

    1 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPCS.  Index: 1


 Timer tables at 086D8F80

 ItTmr        Pass     Slot      Delta        Max   PopCount Array@
 086D8F80       64       62        100      12800       8253 086D9000

   Global TQE Queue for Slot 63:

     Tqe      Tid      Ecb      Mod      Parm           Msec TqeFlag  TidFLag

     08EDDD58 08EDDD44 00000000 EZBIFIU2 08EDDD40        100 00       20

   1 TQE(s) for slot 63 with 0 msec timer offset

 ItTmr        Pass     Slot      Delta        Max   PopCount Array@
 086D8FA0        6       58       1000     128000        825 086D9400

 ItTmr        Pass     Slot      Delta        Max   PopCount Array@
 086D8FC0        0       83      10000    1280000         82 086D9800

  Global TQE Queue for Slot 109:

  Tqe      Tid      Ecb      Ascb     Asid           Msec TqeFlags TidFLags

  16DEF310 166E8D68 166E8CE8 00F8E500 001A0000     300000 40       20000000
```

```
1 TQE(s) for slot 109 with 138000 msec timer offset

   Global TQE Queue for Slot 122:

     Tqe      Tid      Ecb      Mod      Parm           Msec TqeFlag  TidFLag

     086C9020 7F4CEBD0 7F4CEBCC 00000000 00000000    1200000 44       20

   1 TQE(s) for slot 122 with 128000 msec timer offset

 ItTmr        Pass     Slot      Delta        Max   PopCount Array@
 086D8FE0        0        9     100000 4294967295          8 086D9C00

2 TQE(s) were found

No cancelled TQE(s) were found

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS TRACE

Use this subcommand to display information about the component trace, the real-time TCP/IP network monitoring NMI, and the real-time application-controlled TCP/IP trace NMI functions.

Syntax

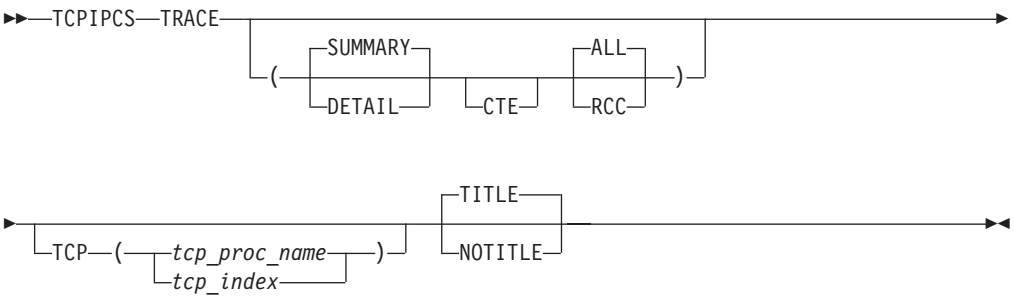

Parameters

SUMMARY
Displays summary information for each function. SUMMARY is the default value. When used with the RCC parameter, the output is similar to the DISPLAY TCPIP,,TRACE,DETAIL command. For more information about this command and its output, see DISPLAY TCPIP,,TRACE command in z/OS Communications Server: IP System Administrator's Commands.

DETAIL
When used with the ALL parameter, lists detailed information about each function in addition to the SUMMARY information. When used with the RCC parameter, displays detailed information about the function.

CTE
Formats each CTE header in each CTRACE buffer. DETAIL must also be specified.

ALL
Formats all the information for all functions.

RCC
Formats only information about the real-time application-controlled TCP/IP trace network management interface (NMI). For more information about this

NMI, see Real-time application-controlled TCP/IP trace NMI in z/OS Communications Server: IP Programmer's Guide and Reference.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {ALL, RCC}, all of the keywords that you specify are used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.
- If you specify CTE, then DETAIL must also be specified.

Sample output of the TCPIPCS TRACE subcommand

The following is sample output of the TCPIPCS TRACE subcommand:

```
TCPIPCS TRACE
Dataset: IPCS.R8A0723.RASDUMP2
Title:   EZRPE005


The address of the TSAB is: 09C445D0

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

09C44610  1 TCPCS     V1R5    093C1000 093C10C8 0029 9FFF7E7F  Active
09C44690  2 TCPCS2    V1R5    00000000 00000000 002A 00000000  Down Stopping

   2 defined TCP/IP(s) were found
   1 active  TCP/IP(s) were found

   2 TCP/IP(s) for CS     V1R5  found

================================================================================

Analysis of Tcp/Ip for TCPCS.  Index: 1

Parmlib Member for SYSTCPIP Trace: CTIEZB00
Parmlib Member for SYSTCPIS Trace: CTIIDS00

Trace Control Area
 TCA: 092BD410
 +0000  TCAACRONYM.......... TCA
 +0006  TCAVERSION.......... 0006
 +0008  TCASIZE............. 0000CBD0
 +000C  TCAFTBE............. 092BD7E0
 +0010  TCACURTBE........... 092BE5C8
 +0014  TCACURENT........... 0059C4C0
 +0018  TCATABSZ............ 01000000
 +001C  TCANUMBF............ 00000100
 +0020  TCABUFSZ............ 00010000
 +0024  TCAMXDAT............ 00003800
 +0028  TCAALET............. 01FF000C
 +002C  TCARCNT............. 00004103
 +0030  TCAECNT............. 00004103
 +0034  TCALCNT............. 00000000
 +0038  TCALTOD............. 00000000  00000000
 +0040  TCACOMP............. 00000000
 +0044  TCAFLAG............. 03200A80
 +0048  TCAXWRTSEQ.......... 00000059
 +004C  TCACTSSWTKN......... 00000000  00000000
 +0054  TCAACNT............. 0000

         -- Array elements --
 +0058  TCAFILTER_ASID...... 0000
```

```
+005A  TCAFILTER_ASID...... 0000
+005C  TCAFILTER_ASID...... 0000
+005E  TCAFILTER_ASID...... 0000
+0060  TCAFILTER_ASID...... 0000
+0062  TCAFILTER_ASID...... 0000
+0064  TCAFILTER_ASID...... 0000
+0066  TCAFILTER_ASID...... 0000
+0068  TCAFILTER_ASID...... 0000
+006A  TCAFILTER_ASID...... 0000
+006C  TCAFILTER_ASID...... 0000
+006E  TCAFILTER_ASID...... 0000
+0070  TCAFILTER_ASID...... 0000
+0072  TCAFILTER_ASID...... 0000
+0074  TCAFILTER_ASID...... 0000
+0076  TCAFILTER_ASID...... 0000
       -- End of array   --

+0078  TCAUCNT............. 0000
       -- Array elements --
+007C  TCAFILTER_USERID.... ........
+0084  TCAFILTER_USERID.... ........
+008C  TCAFILTER_USERID.... ........
+0094  TCAFILTER_USERID.... ........
+009C  TCAFILTER_USERID.... ........
+00A4  TCAFILTER_USERID.... ........
+00AC  TCAFILTER_USERID.... ........
+00B4  TCAFILTER_USERID.... ........
+00BC  TCAFILTER_USERID.... ........
+00C4  TCAFILTER_USERID.... ........
+00CC  TCAFILTER_USERID.... ........
+00D4  TCAFILTER_USERID.... ........
+00DC  TCAFILTER_USERID.... ........
+00E4  TCAFILTER_USERID.... ........
+00EC  TCAFILTER_USERID.... ........
+00F4  TCAFILTER_USERID.... ........
       -- End of array   --

+0100  TCAPCNT............. 00000000

       -- Array elements --
+0104  TCAFILTER_PORT...... 0000
+0106  TCAFILTER_PORT...... 0000
+0108  TCAFILTER_PORT...... 0000
+010A  TCAFILTER_PORT...... 0000
+010C  TCAFILTER_PORT...... 0000
+010E  TCAFILTER_PORT...... 0000
+0110  TCAFILTER_PORT...... 0000
+0112  TCAFILTER_PORT...... 0000
+0114  TCAFILTER_PORT...... 0000
+0116  TCAFILTER_PORT...... 0000
+0118  TCAFILTER_PORT...... 0000
+011A  TCAFILTER_PORT...... 0000
+011C  TCAFILTER_PORT...... 0000
+011E  TCAFILTER_PORT...... 0000
+0120  TCAFILTER_PORT...... 0000
+0122  TCAFILTER_PORT...... 0000
       -- End of array   --

+0124  TCAICNT............. 00000000
       -- Array elements --
+0128  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
+0138  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
+0148  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
+0158  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
+0168  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
+0178  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
+0188  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
```

```
 +0198  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +01A8  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +01B8  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +01C8  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +01D8  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +01E8  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +01F8  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +0208  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +0218  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +0228  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +0238  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +0248  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +0258  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +0268  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +0278  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +0288  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +0298  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +02A8  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +02B8  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +02C8  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +02D8  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +02E8  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +02F8  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
 +0308  TCAFILTER_IPADDRESS. 00000000  00000000  00000000  00000000
 +0318  TCAFILTER_IPSUBMASK. 00000000  00000000  00000000  00000000
         -- End of array   --

 +0330  TCAWRAPTIME......... B66479D0  DE1B3402
 +0338  TCAWRAPCOUNT........ 00000001
 +033C  TCAXWTRCNT.......... 00000000
 +0340  TCACURCUR........... 092BFFE0
 +0344  TCANXTCUR........... 01004E80
 +0348  TCADATSZ............ 02000000
 +034C  TCADATBF............ 00000200
 +0350  TCAWCONT............ 00000068
 +0354  TCATRCNT............ 00000068
 +0358  TCALSCNT............ 00000000
 +035C  TCASEQXWRT.......... 00000001
 +0360  TCAPTSSWTKN......... 02895060  00000001
 +0368  TCAISTBE............ 092C4FE0
 +036C  TCAISNRTBE.......... 00000200
 +0370  TCAISBUFSZ.......... 00010000
 +0374  TCAISTBLSZ.......... 02000000
 +0378  TCAISTRCNT.......... 00000000
 +037C  TCAISWRCNT.......... 00000000
 +0380  TCAISLSCNT.......... 00000000
 +0384  TCAISRQCNT.......... 00000000
 +0388  TCAISXWSEQ.......... 00000001
 +0390  TCAISCDS............ 092C4FE0  03001000
 +0390  TCAISCDTBE.......... 092C4FE0
 +0394  TCAISCDBUF.......... 03001000
 +0398  TCAISXWTKN.......... 00000000  00000000
 +03A0  TCAISWRTIM.......... 00000000  00000000
 +03A8  TCAISLSTIM.......... 00000000  00000000
 +03B0  TCADUMPSZ........... 00000000
 +03B4  TCADUMPDS........... 00000000
 +03B8  TCADUMPOF........... 00000000
 +03C0  TCADUMPTOD.......... 00000000  00000000
Event Trace Statistics for SYSTCPIP
  Size of the Trace Control Area . . . . 52176
  Size of the trace buffer . . . . . . . 16384K
  Size of a trace segment. . . . . . . . 64K
  Number of trace segments . . . . . . . 256
  Maximum trace record size. . . . . . . 14,336
  Number of trace records requested. . . 16,643
  Number of trace records recorded . . . 16,643
  Number of trace segments filled. . . . 89
```

```
  Average records per segment. . . . . 187
  Average records per table. . . . . 47,872
  Trace status . . . . . Active
  XWriter status . . . . . Disconnected
  Number of buffers written. . . . . 0
  Lost record count. . . . . 0
  Lost record time . . . . . 1900/01/01 00:00:00.000000
  Trace table wrap count . . . . . 1
  Trace table wrap time. . . . . 2001/09/05 12:41:47.461043
  Average records per wrap . . . . . 16,643

Data Trace Statistics for SYSTCPDA
  Size of the trace buffer . . . . . 32768K
  Size of a trace segment. . . . . 64K
  Number of trace segments . . . . . 512
  Number of trace records requested. . . 104
  Number of trace records recorded . . . 104
  Number of trace segments filled. . . . 1
  Trace status . . . . . Active
  XWriter status . . . . . Connected
  Number of lost records . . . . . 0

IDSTRACE Statistics for SYSTCPIS
  Size of the trace buffer . . . . . 32768K
  Size of a trace segment. . . . . 64K
  Number of trace segments . . . . . 512
  Number of trace records requested. . . 0
  Number of trace records recorded . . . 0
  Number of trace segments filled. . . . 1
  Trace status . . . . . Active
  XWriter status . . . . . Disconnected
  Number of lost records . . . . . 0
  Lost record time . . . . . 1900/01/01 00:00:00.000000
  Trace table wrap count . . . . . 0
  Trace table wrap time. . . . . 1900/01/01 00:00:00.000000

Tseb_Trace_Opts: 9FFF7E7F
Options: Init Opcmds Opmsgs Socket AFP XCF Access PFS API
         Engine Streams Queue RAW UDP TCP ICMP ARP CLAW LCS
         Internet Message WorkUnit Config SNMP IOCTL FireWall
         VtamData TelnVtam Telnet Vtam

 256 SYSTCPIP Trace Buffer Elements were found
   0 SYSTCPIP Trace Buffer Elements were formatted
 512 SYSTCPDA Trace Buffer Elements were found
   0 SYSTCPDA Trace Buffer Elements were formatted
 512 SYSTCPIS Trace Buffer Elements were found
   0 SYSTCPIS Trace Buffer Elements were formatted


Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS TREE

Use this subcommand to display the structure of TCP/IP Patricia trees.

Syntax

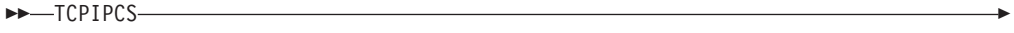

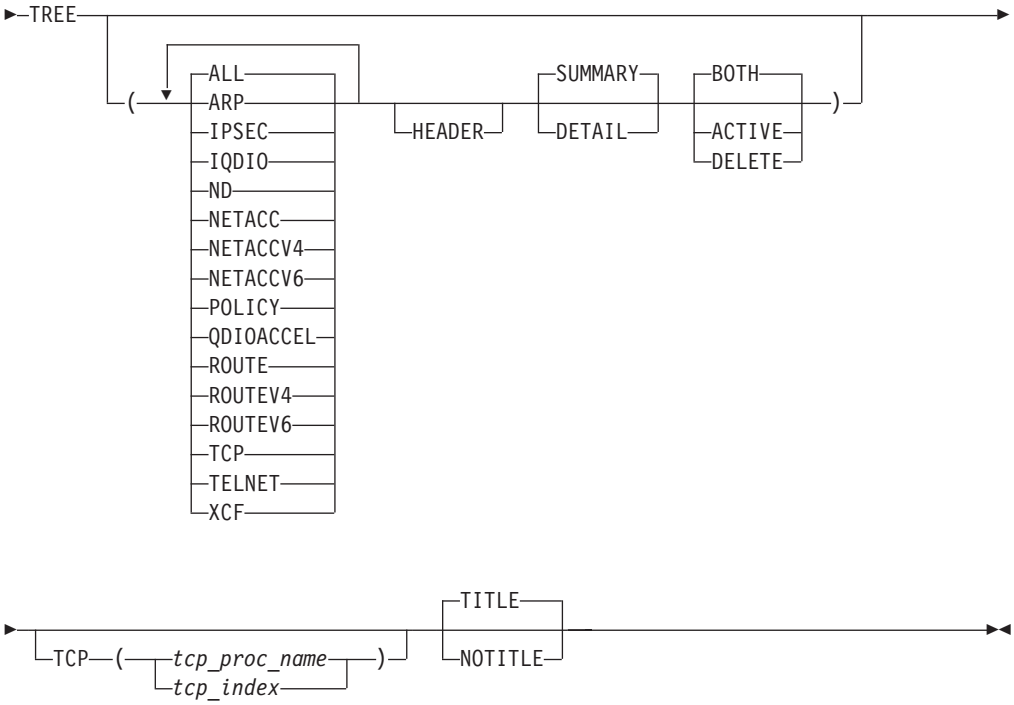

Parameters

ALL
Display structure of all TCP/IP trees. ALL is the default.

ARP
Display only structure of ARP trees.

IPSEC
Display only structure of IP security trees.

IQDIO
Display only structure of iQDIO and QDIOACCEL trees.

ND Display only structure of Neighbor Discovery trees.

NETACC
Display only structure of NetAccess trees.

NETACCV4
Display only structure of IPv4 NetAccess trees.

NETACCV6
Display only structure of IPv6 NetAccess trees.

POLICY
Display only structure of Service Policy trees.

QDIOACCEL
Display only structure of iQDIO and QDIOACCEL trees.

ROUTE
Display only structure of both IPv4 and IPv6 route trees.

ROUTEV4
Display only structure of IPv4 route trees.

ROUTEV6
Display only structure of IPv6 route trees.

TCP
Display only structure of TCP trees.

TELNET
Display only structure of Telnet trees.

XCF
Display only structure of XCF trees.

HEADER
Display tree header information. Not displayed by default.

SUMMARY
Display the addresses of the control blocks and other data in trees. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows the search key values.

BOTH
Display both active and logically deleted tree nodes. BOTH is the default.

ACTIVE
Display only active tree nodes.

DELETE
Display only logically deleted tree nodes

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {ALL, ARP, IPSEC, IQDIO, ND, NETACC, NETACCV4, NETACCV6, POLICY, QDIOACCEL, ROUTE, ROUTEV4, ROUTEV6, TCP, TELNET, XCF}, all of the keywords that you specify are used.
- If you specify multiple keywords from the set {BOTH, ACTIVE, DELETE}, only the last one is used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS TREE subcommand

The following is sample output of the TCPIPCS TREE subcommand:

```
TCPIP Tree Analysis

IPv4 NetAccess Search Tree

Node@     Bit Parent   LChild   RChild   Key      Element
2B42D678  255 00000000 7F042D90 7F04D6F0
7F04D230    2 7F0D1010 7F0D1010 7F04D230 7F04D490 2B1F4898
7F0D1010    3 7F04D570 7F04D6F0 7F04D230 7F04D190 2B1F48F8
7F04D570    4 7F04CD90 7F0D1010 7F04D570 7F04D610 2B1F4838
7F04CC50    9 7F04C950 7F04C950 7F04CC50 7F04CCF0 2B1F47D8
7F04C810    9 7F04C510 7F04C510 7F04C810 7F04C8B0 2B2064B8
7F04C250    9 7F04C510 7F0D4010 7F04C250 7F04C470 2B206578
7F04C510   10 7F0D4010 7F04C250 7F04C810 7F04C770 2B206518
7F04BD90    9 7F04BB10 7F04BB10 7F04BD90 7F04BFB0 2B206638
7F04B850    9 7F04BB10 7F04B5D0 7F04B850 7F04BA70 2B2066F8
7F04BB10   10 7F0D4010 7F04B850 7F04BD90 7F04BCF0 2B206698
7F042D90   32 2B42D678 2B42D678 7F042D90 7F042CF0 2B1F4658

11 elements in IPv4 NetAccess Search Tree
```

```
IPv4 NetAccess Update Tree

Node@     Bit Parent   LChild   RChild   Key      Element
2B21E818  255 00000000 7F042E10 7F04D670
7F04D2B0    2 7F04D030 7F04D030 7F04D2B0 7F04D430 2B1F4898
7F04D030    3 7F04D4F0 7F04D670 7F04D2B0 7F04D130 2B1F48F8
7F04D4F0    4 7F04CE10 7F04D030 7F04D4F0 7F04D5B0 2B1F4838
7F04CE90    9 7F04C9D0 7F04C9D0 7F04CE90 7F04CC90 2B1F47D8
7F04CA50    9 7F04C590 7F04C590 7F04CA50 7F04C850 2B2064B8
7F04C2D0    9 7F04C590 7F04C050 7F04C2D0 7F04C410 2B206578
7F04C590   10 7F04C050 7F04C2D0 7F04CA50 7F04C710 2B206518
7F04BE10    9 7F04BB90 7F04BB90 7F04BE10 7F04BF50 2B206638
7F04B8D0    9 7F04BB90 7F04B650 7F04B8D0 7F04BA10 2B2066F8
7F04BB90   10 7F04C050 7F04B8D0 7F04BE10 7F04BC90 2B206698
7F042E10   32 2B21E818 2B21E818 7F042E10 7F042C90 2B1F4658

11 elements in IPv4 NetAccess Update Tree

IPv6 NetAccess Search Tree

Node@     Bit Parent   LChild   RChild   Key      Element
2B2180B8  255 00000000 7F04D830 7F0A0010
7F085010    2 7F083010 7F083010 7F085010 7F050B70 2B1F1658
7F079350    1 7F080010 7F0A0010 7F079350 7F050930 2B1F1598
7F080010    2 7F083010 7F079350 7F080010 7F0509F0 2B1F1538
7F083010    3 7F0C4010 7F080010 7F085010 7F050AB0 2B1F16B8
7F0C4010    4 7F0506D0 7F083010 7F0C4010 7F050C30 2B1F15F8
7F050510    2 7F0506D0 7F0506D0 7F050510 7F050630 2B1F1478
7F0506D0  114 7F0A3010 7F050510 7F0C4010 7F050470 2B1F14D8
7F0A3010  115 7F050290 7F0506D0 7F0A3010 7F050CF0 2B1F1418
7F0503D0   65 7F0664D0 7F0664D0 7F0503D0 7F050170 2B1F1898
7F0661D0   65 7F0664D0 7F0C9010 7F0661D0 7F0663B0 2B1F1958
7F0664D0   66 7F0C9010 7F0661D0 7F0503D0 7F0666B0 2B1F18F8
7F04FDD0   65 7F04FB10 7F04FB10 7F04FDD0 7F04FFB0 2B1F40B8
7F04F790   65 7F04FB10 7F04F510 7F04F790 7F04F9F0 2B1F4178
7F04FB10   66 7F0C9010 7F04F790 7F04FDD0 7F04FCB0 2B1F4118
7F04D830  128 2B2180B8 2B2180B8 7F04D830 7F04D790 2B1F12F8

15 elements in IPv6 NetAccess Search Tree

IPv6 NetAccess Update Tree

Node@     Bit Parent   LChild   RChild   Key      Element
2B42D6D8  255 00000000 7F04D7F0 7F0B2010
7F0685D0    2 7F08D010 7F08D010 7F0685D0 7F050B10 2B1F1658
7F079250    1 7F079310 7F0B2010 7F079250 7F0508D0 2B1F1598
7F079310    2 7F08D010 7F079250 7F079310 7F050990 2B1F1538
7F08D010    3 7F0C2010 7F079310 7F0685D0 7F050A50 2B1F16B8
7F0C2010    4 7F050690 7F08D010 7F0C2010 7F050BD0 2B1F15F8
7F050590    2 7F050690 7F050690 7F050590 7F0505D0 2B1F1478
7F050690  114 7F0A8010 7F050590 7F0C2010 7F050410 2B1F14D8
7F0A8010  115 7F050310 7F050690 7F0A8010 7F050C90 2B1F1418
7F050010   65 7F066550 7F066550 7F050010 7F050110 2B1F1898
7F066250   65 7F066550 7F0C7010 7F066250 7F066350 2B1F1958
7F066550   66 7F0C7010 7F066250 7F050010 7F066650 2B1F18F8
7F04FE50   65 7F04FB90 7F04FB90 7F04FE50 7F04FF50 2B1F40B8
7F04F810   65 7F04FB90 7F04F590 7F04F810 7F04F990 2B1F4178
7F04FB90   66 7F0C7010 7F04F810 7F04FE50 7F04FC50 2B1F4118
7F04D7F0  128 2B42D6D8 2B42D6D8 7F04D7F0 7F04D730 2B1F12F8

15 elements in IPv6 NetAccess Update Tree
```

TCPIPCS TSDB

Use this subcommand to display the TSDB server data block.

Syntax

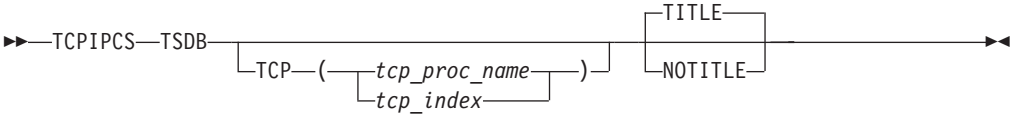

Parameters

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS TSDB subcommand

The following is sample output of the TCPIPCS TSDB subcommand:

```
TCPIPCS TSDB
 Dataset: IPCS.MV21381.DUMPA
 Title:   SLIP DUMP ID=TC


 The address of the TSAB is: 13391BC0

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 13391C00  1 TCPSVT    V2R10   1323B000 1323B0C8 07DE 04041405  Active
 13391C80  2 TCPSVT2   V2R10   00000000 00000000 07E8 00000000  Down Stopping
 13391D00  3 TCPSVT1   V2R10   12FC3000 12FC30C8 0080 94FF755F  Active
 13391D80  4 TCPSVT3   V2R10   00000000 00000000 0059 00000000  Down Stopping

    4 defined TCP/IP(s) were found
    2 active  TCP/IP(s) were found

    4 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPSVT.  Index: 1

 TSDB control block summary

  TSDB: 1323B000
  +0000  TSDB_ACRONYM............. TSDB
  +0004  TSDB_LENGTH.............. 00C8
  +0006  TSDB_VERSION............. 0003
  +0008  TSDB_STATE............... 0015
  +000A  TSDB_ASID................ 07DE

          -- Array elements --
  +0010  TSDB_MT.................. 11A7E870
  +0014  TSDB_MT.................. 962F5E00
  ....
  +0060  TSDB_CTRACE_PARMLIB_NAME. CTIEZB02
  +006C  TSDB_SMCA................ 00000000
  +0070  TSDB_TSRMT............... 00000000
  +0074  TSDB_FLAGS............... 00000000
  +0078  TSDB_CONFIG_PORT......... 00000401
  +007C  TSDB_OSASF_PORT.......... FFFFFFFF
  +0080  TSDB_EZBITMSN@........... 91A8BF90
  +0084  TSDB_TERMINATING_ECB..... 807EC758
  +0088  TSDB_DUAF................ 00000000
  +008C  TSDB_TSCA................ 13236A58
  +0090  TSDB_SOCIFPTR............ 91BC3E78
  +0094  TSDB_SOMIFPTR............ 91BCA050
  +0098  TSDB_RXGLUPTR............ 91BF6308
  +00A8  TSDB_LEPARMS............. 14B01BBA
```

```
 +00AC  TSDB_OE_AS_STOKEN........ 00000038  00000001
 +00B4  TSDB_SOMT2............... 91C3FE60

Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS TSDX

Use this subcommand to display the TSDX server data extension.

Syntax

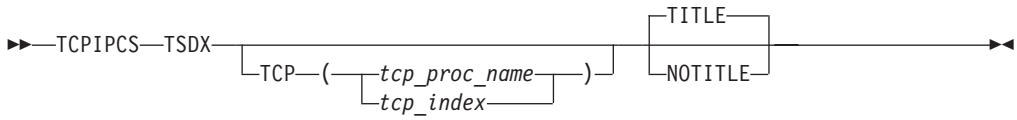

Parameters

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS TSDX subcommand

The following is sample output of the TCPIPCS TSDX subcommand:

```
TCPIPCS TSDX
Dataset: IPCS.MV21381.DUMPA
Title:   SLIP DUMP ID=TC


The address of the TSAB is: 13391BC0

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

13391C00  1 TCPSVT    V2R10   1323B000 1323B0C8 07DE 04041405  Active
13391C80  2 TCPSVT2   V2R10   00000000 00000000 07E8 00000000  Down Stopping
13391D00  3 TCPSVT1   V2R10   12FC3000 12FC30C8 0080 94FF755F  Active
13391D80  4 TCPSVT3   V2R10   00000000 00000000 0059 00000000  Down Stopping

   4 defined TCP/IP(s) were found
   2 active  TCP/IP(s) were found

   4 TCP/IP(s) for CS     V2R10 found

================================================================================

Analysis of Tcp/Ip for TCPSVT.  Index: 1

TSDX control block summary

 TSDX: 1323B0C8
 +0000  TSDX_ACRONYM............... TSDX
 +0004  TSDX_LENGTH................ 0300
 +0006  TSDX_VERSION............... 0003
 +0008  TSDX_FLAGS................. 60000001
 +000C  TSDX_ASCB.................. 00F7C280
 +0010  TSDX_PROCNAME.............. TCPSVT
 +0018  TSDX_CART.................. 00000000  00000000
 +0020  TSDX_CONSID................ 00000001
 +0024  TSDX_TCB................... 007EC9A8
 +0028  TSDX_TCB_TOKEN............. 00001F78  00000008  00000003  007EC9A8
 +0038  TSDX_TCPIP_DS_ALET......... 01FF0011
 +003C  TSDX_TCPIP_DS_ADDR......... 00001000
 +0040  TSDX_TCPIP_DS_END.......... 19001000
 +0044  TSDX_ET_TOKEN.............. 7FFD9D10
...
```

```
 +026C  TSDX_CSMSTATAREA........... 141C7A88
 +0270  TSDX_CSMDUMPINFO........... 141C7A90
 +0288  TSDX_AUTOLOG_TASK_ECB...... 807EC758
 +028C  TSDX_AUTOLOG_CB............ 1333C0A8
 +0290  TSDX_SASTRT_ECB............ 807EC758
 +0294  TSDX_XFCVT................. 13096410
 +0298  TSDX_XCFLOCK............... 00000000  00000000  00000000  D7D60901

Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS TSEB

Use this subcommand to display the TSEB server anchor block.

Syntax

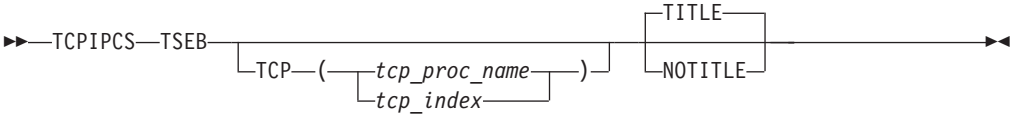

Parameters

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Sample output of the TCPIPCS TSEB subcommand

The following is sample output of the TCPIPCS TSEB subcommand:

```
TCPIPCS TSEB
 Dataset: IPCS.MV21381.DUMPA
 Title:   SLIP DUMP ID=TC


 The address of the TSAB is: 13391BC0

 Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

 13391C00  1 TCPSVT    V2R10   1323B000 1323B0C8 07DE 04041405  Active
 13391C80  2 TCPSVT2   V2R10   00000000 00000000 07E8 00000000  Down Stopping
 13391D00  3 TCPSVT1   V2R10   12FC3000 12FC30C8 0080 94FF755F  Active
 13391D80  4 TCPSVT3   V2R10   00000000 00000000 0059 00000000  Down Stopping

    4 defined TCP/IP(s) were found
    2 active  TCP/IP(s) were found

    4 TCP/IP(s) for CS     V2R10 found

 ================================================================================

 Analysis of Tcp/Ip for TCPSVT.  Index: 1

 TSEB control block summary

  TSEB: 13391C00
  +0000  TSEB_ACRONYM.......... TSEB
  +0004  TSEB_LENGTH........... 0080
  +0006  TSEB_VERSION.......... 0003
  +0008  TSEB_FLAGS............ 82000000
  +0008  TSEB_STATUS........... 82
  +000C  TSEB_REQUESTORS....... 00000000
  +0010  TSEB_TCPIP_NAME....... TCPSVT
  +0018  TSEB_SI............... 01
  +0019  TSEB_IID.............. 04
  +001A  TSEB_TCPIP_VERSION.... 0510
```

```
  +001C  TSEB_TSDB.............. 1323B000
  +0020  TSEB_LX................ 00002E00
  +0024  TSEB_TCA............... 11469E50
  +0028  TSEB_TRACE_OPTS....... 04041405
  +002C  TSEB_TRACE_OPT2....... 00000000
  +0034  TSEB_SCHEDULED_EVENTS. 00000000
  +0038  TSEB_ASID.............. 07DE
  +003C  TSEB_LPA_SADDR........ 11A719E0
  +0040  TSEB_LPA_EADDR........ 11C62FFF
  +0044  TSEB_QDIO_BGRP_Q@..... 1320A648
  +0048  TSEB_EZBITDCR......... 9320ACF8
  +004C  TSEB_ITCVT............ 1323B3C8
  +0050  TSEB_BGRP_Q@.......... 1320A608
  +0054  TSEB_DUAF............. 130A4010
  +0059  TSEB_TOKENID.......... 000015
  +005C  TSEB_TCMTPTR.......... 132072F0
  +0060  TSEB_EZBITCOM_LEN..... 00007D28
  +0064  TSEB_CS390_VERSION.... 020A
  +0068  TSEB_CSMFREE.......... 1320A548
  +006C  TSEB_TCPIP_STOKEN..... 00001F78  00000008
  +0074  TSEB_CSMPACK.......... 1320A588
  +0078  TSEB_SOCA............. 140BAEB8
  +007C  TSEB_CSMPACKQDIO...... 1320A5C8

 Analysis of Tcp/Ip for TCPSVT completed
```

TCPIPCS TTLS

Display information about Application Transparent Transport Layer Security (AT-TLS), AT-TLS groups, and AT-TLS connections.

Syntax

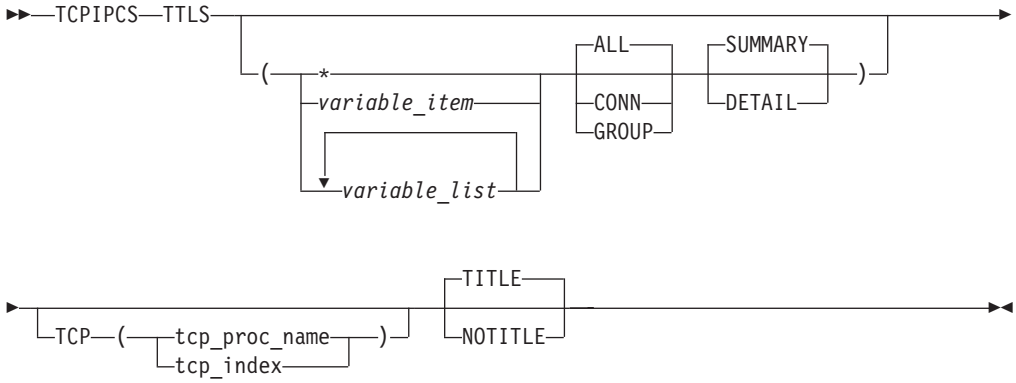

Parameters

If no parameters are specified, both AT-TLS connections and AT-TLS groups are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
From 1 - 32 of the following variable parameters can be repeated, each separated by a blank space, within parentheses:

Variable parameters are:

TCB_address
Displays AT-TLS information for the connection with this address. An address is specified as 1–8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

group_address
Displays information for the AT-TLS group with this address. An address is specified as 1–8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

connection_id
Displays AT-TLS information for the connection with this connection ID. An ID is specified as 1–8 hexadecimal digits.

group_id
Displays information for the AT-TLS group with this group ID. An ID is specified as 1–8 hexadecimal digits.

In addition to the variable parameters described above, the following keyword parameters can be specified:

CONN
Display only information for AT-TLS connections.

GROUP
Display only information for AT-TLS groups.

ALL
Display information for both AT-TLS connections and groups. ALL is the default.

SUMMARY
Displays the addresses of the control blocks and other data in tables. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows the contents of the control blocks.

TCP, TITLE, NOTITLE
See “Parameters” on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {CONN, GROUP, ALL}, only the last one is used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS TTLS subcommand

The following is sample output of the TCPIPCS TTLS subcommand:

```
 TCPIPCS TTLS(* )
.....
 =
 Analysis of Tcp/Ip for TCPCS4.  Index: 2
 TCP/IP Analysis
 TCPIP Main TTLS Control Block (EZBZTTLS)
  EZBZTTLS: 7F1E10F0
```

```
 +0000  TTLS_ACRONYM..... EZBZTTLS
 +0008  TTLS_PART_LOCK..... 00000000  00000000  00000000  D227030A
 +0008  LOCK_CDS..... 00000000  00000000
 +0008  LOCK_SUSPENDED_GLOBAL. 00000000
 +000C  LOCK_HOLDER..... 00000000
 +0010  LOCK_SUSPENDED_LOCAL.. 00000000
 +0014  LOCK_INFO..... D227030A
 +0014  LOCK_INIT..... D227
 +0014  LOCK_INIT1..... D2
 +0015  LOCK_INIT2..... 27
 +0016  LOCK_CLASS..... 03
 +0017  LOCK_LEVEL..... 0A
 +0018  TTLS_FLAG1..... E8000000
 +001C  TTLS_PIPIADD_CNT..... 00
 +001D  TTLS_GRPCNT..... 0A
 +001E  TTLS_TCB_CMPC_OFF..... 11
 +001F  TTLS_ABEND_COUNT..... 00
 +0020  TTLS_1STABEND..... 00000000
 +0024  TTLS_TCBPTR..... 007D61C8
 +0028  TTLS_RESMGR_TOKEN..... 00000669  00000000
 +0030  TTLS_INBNDPART@..... 7E447310
 +0034  TTLS_OUTBNDPART@..... 7E447190
 +0038  TTLS_PCT_STATE..... 0000000F
 +003C  TTLS_PCT_INSTANCEID... 4E42929D
 +0040  TTLS_WORKQ..... 00000000  00000000
 +0040  ITLFPUBLIC..... 00000000
 +0040  ITLFHEAD..... 00000000
 +0044  ITLFPRIVATE..... 00000000
 +0044  ITLFTAIL..... 00000000
 +0048  TTLS_TERM_ECB..... 807D6120
 +004C  TTLS_INIT_ECB..... 00000000
 +0050  TTLS_EOT_ECB..... 007FF180
 +0054  TTLS_WORKQ_ECB..... 807D6120
 +0058  TTLS_CLEANUP_TIMER.... 00000000  00000000  00000000  00000000
 +0058  TID_EYE..... 000000
 +005B  TID_BITS..... 00
 +005C  TID_MSEC..... 00000000
 +0060  TID_CDS..... 00000000  00000000
 +0060  TID_FLAGS..... 00000000
 +0064  TID_TQE..... 00000000
 +0068  TTLS_MODLIST..... 7F1E11B4
 +006C  TTLS_GROUPNUM..... 00000069
 +0070  TTLS_TGRP_HT_TOKEN.... 7F54E070  00000004
 +0078  TTLS_CLEANUP_ECB..... 00000000
 +007C  TTLS_MAX_SRBS..... 00000005
 +0080  TTLS_CURR_SRBS..... 00000000
 +0084  TTLS_WE_CNT..... 00000000
 +0088  TTLS_PIPI_ECB..... 807D6120
 +008C  TTLS_PIPI_POOLPTR..... 7F5439B0
 +0090  TTLS_PIPI_SUSPQ..... 00000000  00000000
 +0090  ITLFPUBLIC..... 00000000
 +0090  ITLFHEAD..... 00000000
 +0094  ITLFPRIVATE..... 00000000
 +0094  ITLFTAIL..... 00000000
 +0098  TTLS_ENVNUM..... 0000002C
 +009C  TTLS_GLBLTHD..... 00000028
 +00A0  TTLS_SECOND_HT_TOKEN.. 7F58B250  00000005
 +00A8  TTLS_MSGQ_CNT..... 00000002
 +00AC  TTLS_LATCH_SET_COUNT.. 00000006
 +00B0  TTLS_LATCH_QUEUE..... 7DE0C760  000000D5
 0 TLMST Work Requests Formatted

TTLS Secondary Map hashtable entries
Pri_TCB@ PID      Local_IP..Remote_IP
7E7A1F10 0200006F ::ffff:10.81.4.4..::ffff:10.81.4.4
TCB@     ConnID   TLSX@    Proto    Cipher Jobname  UserID   Cert@    CertId
7E7A1F10 0000015D 7E308130 TLSV1.2  C02B   USER11   USER1    7E328F18 ........
  LocalSocket:  10.81.4.4..1046
```

```
   RemoteSocket: 10.81.4.4..50500
   Tcb_tcp_state: Established
   Tcb_TtlsFlags:Ttls_Gate Ttls_Enabled Ttls_Started
   TLSX_Flags_B1:  v3Header Fast_Local
   TLSX_Flags_B2:  Check_Reset Data_Seen
   TLSX_Flags_B3:
   TLSX_Flags_B4:  LookUp_Done
   TLSX_Flags_B5:
   TTLSRule(7E3B1B90):       ttls_ruleC
   TTLSGroupAction(7E337D10):  group_action2
   TTLSEnvironmentAction(7E3AAF10):  environment_action2
   TTLSConnectionAction(7E3AB610): connection_action2
....
  15 TCBs Found
   3 TCBs Formatted
....

 TTLS Group: group_action1
   Address Group Id    Conns    Tasks Elements Created
  7E305D70       63        2        4        0 2011/08/10 14:15:57
   -----TTLS Environments---------------
    TTLS Environment: environment_action1
      Address   Env Id    Conns    Rcode Env Type Idle Time(Valid for 0 Conns)
     7E0FF9B0       2B        1        0 Master    1900/01/01 00:00:0
    TTLS Environment: environment_action1
      Address   Env Id    Conns    Rcode Env Type Idle Time(Valid for 0 Conns)
     7E1CB2D0       2A        2        0 Process   1900/01/01 00:00:0
   2 TTLS Environments Formatted
   -----TTLS Worker Tasks---------------
    TTLS Worker Task: 7E25DA50
         Ducb FuncCode    Rcode     Busy Idle Time
     14F3B000       15        0        0 2011/08/10 14:16:55
    TTLS Worker Task: 7E262350
         Ducb FuncCode    Rcode     Busy Idle Time
     14F3E000        3        0        0 2011/08/10 14:15:56
    TTLS Worker Task: 7E25D890
         Ducb FuncCode    Rcode     Busy Idle Time
     1527E000        3        0        0 2011/08/10 14:15:56
    TTLS Worker Task: 7E25D190
         Ducb FuncCode    Rcode     Busy Idle Time
     15296000        3        0        0 2011/08/10 14:15:56
  4 TTLS Worker Tasks Formatted
  0 TGRP Work Requests Formatted
   -----TGRP Log Requests---------------
   ***** Following in Reverse Order *****
 TLWE@    Clock           Function          Parm1    Parm2    ExtLen
 7E3D0C90 14:17:15.810123 WRITE_SYSLOG     00000007 7E3D0CC8 0000007B 00000000
  1 TGRP Log Requests Formatted
 ================================================================================

  10 TTLS Group Found
   1 TTLS Group Formatted
```

TCPIPCS UDP

Use this subcommand to display the Master UDP Control Block (MUCB) and any UDP Control Blocks (UCBs) in the UDP hash tables or link list.

Syntax

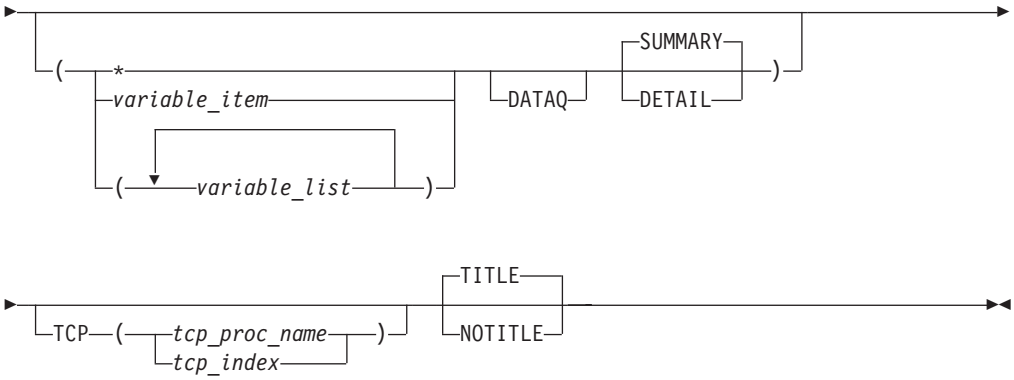

Parameters

If no parameters are specified, all UDP control blocks are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1–32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

jobname
Displays only the UDP control blocks with this job name. The job name can be a TCP/IP application name or a stack name. A job name is 1–8 alphanumeric characters.

UCB_address
Displays only the UDP control block with this address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

connection_id
Displays the UDP control block with this connection ID. A connection ID is specified as 1–8 hexadecimal digits.

In addition to the variable parameters described above, you can specify the following keyword parameters:

DATAQ
Formats UCBs which have data queued on the RECEIVE queue.

SUMMARY
Formats the MUCB contents and lists all the UDPs in one cross-reference table. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of the UCBs.

TCP, TITLE, NOTITLE
See “Parameters” on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS UDP subcommand

The following is sample output of the TCPIPCS UDP subcommand:

```
TCPIPCS UDP
Dataset: IPCS.R8A0723.RASDUMP
Title:   EZRPE005
The address of the TSAB is: 098221F0
Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status
09822230  1 TCPCS     V1R5    08E85000 08E850C8 001E 9FFF7E7F  Active
098222B0  2 TCPCS2    V1R5    08937000 089370C8 01F6 9FFF7E7F  Active
   2 defined TCP/IP(s) were found
   2 active  TCP/IP(s) were found
   2 TCP/IP(s) for CS     V1R5  found

Analysis of Tcp/Ip for TCPCS.  Index: 1
User Datagram Protocol Control Block Summary
 MUCB: 7F50B248
 +0000  UMUCBEYE. MUCB      USTKLNKD. 01        UDRVSTAT. 00
 +0008  UMUCB6FLG.........           00010000
 +0009  U6STKLNKD.........           01
 +000B  U6DRVSTAT.........           00
 +000C  UOPENPRT. 00000000  UFREEPRT. 041C      MCBMUTEX. 00000000
        00000000  00000000  D7D60402
 +0028  UDPCFG... 00000001  0000FFFF  0000FFFF  00000001  80000000
        00000000
 +0040  UDPCFG2.. 00000001  0000FFFF  0000FFFF  00000001  80000000
        00000000
 +0058  UDPMIB... 00000008  0000004B  00000000  0000004D
        USBCAST.. 00000000  USLPBACK. 00000000  USDN
 +0074  USRCVBUF. 0000FFFF  USSNDBUF. 0000FFFF  UFGPRC... 00
        USERIALV. 00000003  USERIAL1. 00000000  ULAS
 +008C  ULASTPRT. 0000      ULASTUCB. 00000000  USERIAL2. 00000000
        UIPWRQ@.. 7F407968  UIPRDQ@.. 7F407928  UIP6
 +00A4  UIP6RDQ@. 7F207928
 +00BC  UDMULTI_NUM........           00000000
 +00C0  UDMUX_TOKEN........           7F407B88  00000003
 +00D0  UDMULTI@. 00000000
 +00D4  UD6MULTI_NUM.......           00000000
 +00D8  UD6MULTI@..........           00000000
 +00DC  UD6MUX_TOKEN.......           7F207B88  00000006
 +00E4  UD6MUX_MULTI_TOKEN.           00000000  00000000

IPv6 Unicast Hash Table
UCB      ResrcID  ResrcNm  TpiState Port IPAddr
7F2FCD00 0000000C TCPCS    WLOUNBND      ::0
1 UCB(s) FOUND
1 UCB(s) FORMATTED

IPv4 Unicast Hash Table
UCB      ResrcID  ResrcNm  TpiState Port IPAddr
7F50C000 00000004 TCPCS    WLOUNBND      0.0.0.0
1 UCB(s) FOUND
1 UCB(s) FORMATTED

IPv6 Multicast Hash Table
0 UCB(s) FOUND
0 UCB(s) FORMATTED

IPv4 Multicast Link List
0 UCB(s) FOUND
0 UCB(s) FORMATTED

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS VMCF

Use this subcommand to display information about VMCF (Virtual Machine Communication Facility) and IUCV (Inter-User Communication Vehicle) users.

Syntax

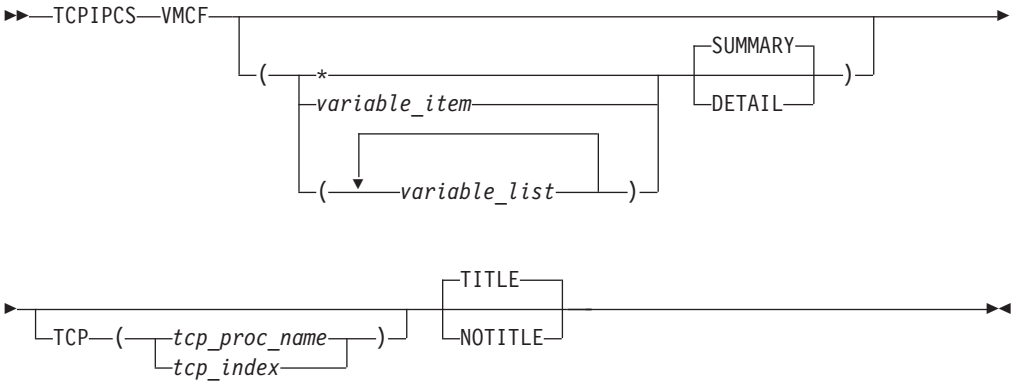

Parameters

If no parameters are specified, all VMCF control blocks are summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

user_id
Displays only the VMCF control block associated with this user ID. A user_id is 1 - 8 alphanumeric characters.

ASCB_address
Displays only the VMCF control blocks associated with this address space control block address. An address is specified as 1-8 hexadecimal digits. An IPCS symbol name can be specified for an address. If an address begins with digit a–f or A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

ASID_number
Displays only the VMCF control blocks associated with this address space identifier. An ASID is specified as one to four hexadecimal digits.

In addition to the variable parameters described above, you can specify the following keyword parameters:

SUMMARY
Formats the VMCF control blocks in one cross-reference table. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of selected VMCF USER control blocks.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rule: If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS VMCF subcommand

The following is sample output of the TCPIPCS VMCF subcommand:

```
TCPIPCS VMCF ((* )  SUMMARY)
Dataset: IPCS.JW11111.DUMPA
Title:   IPCS VMCF DUMP


The address of the TSAB is: 08EBC180

Tseb     SI Procedure Version Tsdb     Tsdx     Asid TraceOpts Status

08EBC1C0  1 TCPCS     V2R10   089DC000 089DC0C8 01F7 9FFFFF7F  Active

   1 defined TCP/IP(s) were found
   1 active  TCP/IP(s) were found

   1 TCP/IP(s) for CS     V2R10 found

================================================================================

Analysis of Tcp/Ip for TCPCS.  Index: 1

TCPIP VMCF Analysis

XINF at 09813000
  VMCF CVT            : 00A44078
  User Array          : 09813090
  Userid Count        : 1
  Userid Array        : 09817050
  Userid              : VMCF
  MSGBUILD            : 89802838
  MVPMSGS             : 8981A290
  Ecb                 : 00000000
  TNF CVT             : 00A63808
  VMCF QD             : 00000000
  VMCF QD Count       : 0
  TNF Manager Area    : 00008FE0
  SMSG Id             : 0

USER at 09813C50
  Userid              : USER18
  Asid                : 005D
  No UserData

Analysis of Tcp/Ip for TCPCS completed
```

TCPIPCS XCF

Use this subcommand to produce a cross-system coupling facility (XCF) analysis report.

Syntax

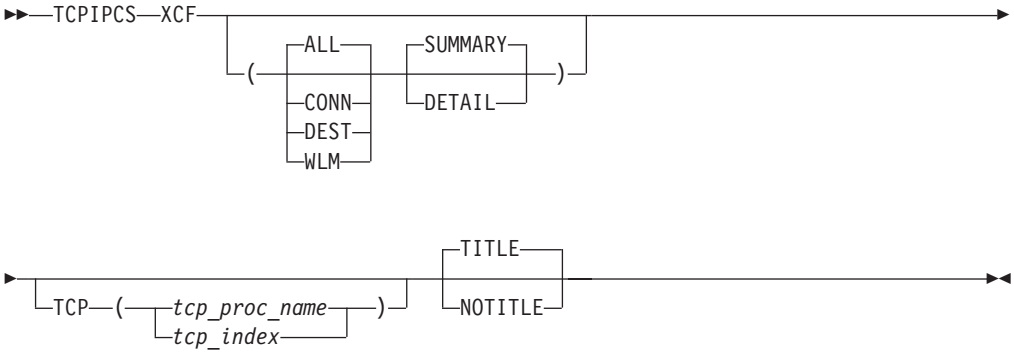

Parameters

If no parameters are specified, the dynamic VIPA hash table and partner tables are summarized.

CONN
Display only connection hash table optional information.

DEST
Display only destination hash table optional information.

WLM
Display only workload manager optional information.

ALL
Display all optional information. ALL is the default.

SUMMARY
Formats the XCF control blocks in one cross-reference table. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL formats the contents of XCF control blocks.

TCP, TITLE, NOTITLE
See "Parameters" on page 177 for a description of these parameters.

Rules:

- If you specify multiple keywords from the set {ALL, CONN, DEST, WLM}, all of the keywords that you specify are used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.

Sample output of the TCPIPCS XCF subcommand

The following is sample output of the TCPIPCS XCF subcommand:

```
TCPIPCS XCF
 Dataset: IPCS.MV20603.DUMPA
 ...
 ----XFCVT information----

 XFCVT@        12CC7410       Member Name RUSSIATCPSVT
 Local PTB     12CC752C       PTB Chain   1276A410
 DVIPAHashT@   13239408       IPHashT@    12A9C010
 ConnRteHashT@ 12A9B010       DPTHashT@   1277D010
 WLMData@      00000000       PolicyPart@ 7F635108
 ================================================================================
 ----DVIPA Hash Table----
 DVIPA Hash Table at 13239408
```

```
Hash table has size 2056 bytes

DVIPA address     197.11.200.2    index 3
         MVSName/TCPName               Status/Rank/@ of BackEle

         RUSSIA/TCPSVT                 Active/255/7C0B2B30
         GERMANY/TCPSVT                Backup/10/7E53C6B0
...
Found 9 entries in the DVIPA Hash Table.
================================================================================
=========================================================
                          Local Partner Table
=========================================================
----Partner Table Control Block----
Partner Table at 12CC752C
NextPtr:    00000000
MVSName:    RUSSIA           CPName:        RUSSIA
TCPname:    TCPSVT           IPTable:       12D6F140
IPCount:    21               IPEntries@:    1322D0E8
...
----Dynamic VIPA Table----
Sending Partner@:  128E1410  GERMANY/TCPSVT

Current Dynamic Home Address: 199.11.87.104
Table Address:  12A98C10     Table Length:  8208
Number of Table Entries:     7
.....................................
 DVIPA entry at 12A98C40
  DVIPA origin: DEFINE       Dist Status: Unknown:0
  DVIPA Flags: MoveImmed
  DVIPA Flag2: ()
  IP address: 197.11.104.10  Mask: 255.255.255.0
...
=========================================================
                          Next Partner Table
=========================================================
----Partner Table Control Block----
Partner Table at 1276A410
NextPtr:    12659410
MVSName:    SPAIN            CPName:        SPAIN
TCPname:    TCPSVT           IPTable:       13BA63A0
...
```

ERRNO command

Use the ERRNO command to search for the name and description of constants used for ERRNO, ErrnoJr, module ID, reason code, and ABEND reason code.

Syntax

Parameters

type
The optional type of value provided:

A Abend code

E Errno

J ErrnoJr

M Module ID

R Reason code (default)

value
The decimal or hexadecimal value to be converted. By default, the value is assumed to be a hexadecimal number. If the value is less than the maximum size for its type, the value is padded on the left with zeros. Choices are:

hhhhhhhh
An address consisting of 1-8 hexadecimal digits ending with a period. The value at that address is interpreted.

hhhhhhhh
An ERRNO, ERRNO junior, reason code, ABEND code, or module ID consisting of 1-8 hexadecimal digits.

hhhhhhhhx
An ERRNO, ERRNO junior, or a module ID consisting of 1-8 hexadecimal digits followed by the letter x.

ddddddddn
An ERRNO, ERRNO junior, or a module ID consisting of 1-8 decimal digits followed by the letter n.

name
The name of a module, an ERRNO, an ErrnoJr, or an ABEND reason code.

Note: If the name is not found, ERRNO attempts to interpret the name as a hexadecimal value.

Sample output of the ERRNO command

This section shows sample outputs of the ERRNO command.

For reason code by hexadecimal value output, code the following:

```
Command ===>  errno r 74be72e9

    ReasonCode: 74BE72E9
      Module: EZBITSTO  ErrnoJr: 29417 JRCMNOCSM
      Description: Cache Manager encountered a CSM storage shortage
```

For reason code by address, where the value at address 07093F98 is 74717273, code the following. Type R (reason code) is the default.

```
    Command ===> errno 7093f98.

     ReasonCode: 74717273
       Module: EZBPFWRT  ErrnoJr: 29299 JRARPSVNOTDEFINED
       Description: The ATMARPSV name specified is not defined
```

For reason code by Errno in decimal, code the following:

```
    Command ===> errno e 129n

    Errno: 00000081(129) : ENOENT
      Description: No such file, directory, or IPC member exists
```

For reason code by ErrnoJr in hexadecimal, code the following:

```
Command ===> errno j 6c

ErrnoJr: 0000006C(108) : JRFILENOTTHERE
  Description: The requested file does not exist
```

For reason code by abend code in decimal, code the following:

```
Command ===>  errno a 9473n

Abend Reason Code: 00002501
  Module: Unknown  Reason: TcpitStorNoCSMstorage
  Description: No CSM storage available
```

For reason code by module ID in hexadecimal, code the following:

```
Command ===> errno m 74be

ModuleId: 74BE(29886) : EZBITSTO EZBTIINI
```

For reason code by module name, code the following:

```
Command ===> errno ezbifinb

ModuleId: 7418(29720) : EZBIFINB EZBTIINI
```

For reason code by ERRNO name, code the following:

```
Command ===> errno ebadf

Errno: 00000071(113) : EBADF
  Description: The file descriptor is incorrect
```

For reason code by ErrnoJr name, code the following:

```
Command ===> errno   jrmaxuids

ErrnoJr: 00000013(19) : JRMAXUIDS
  Description: The maximum number of OpenMVS user IDs is exceeded
```

For reason code by ABEND reason name, code the following:

```
Command ===> errno tcpbadentrycode

Abend Reason Code: 00000401
  Module: Unknown  Reason: TcpBadEntryCode
  Description: Bad Entry code to module
```

ICMPHDR command

This section describes the ICMPHDR command.

Use the ICMPHDR command to display the ICMP header fields.

Syntax

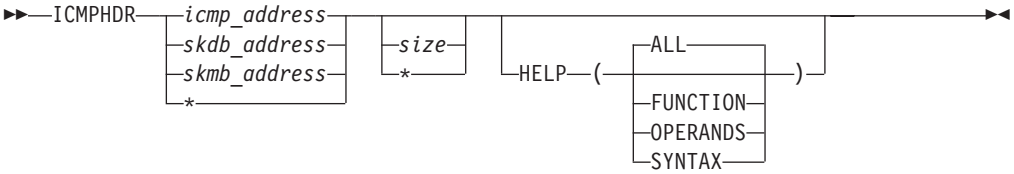

Parameters

***** To omit this positional parameter when using the HELP keyword.

icmp_address
The address of an ICMP header or the symbol for the address.

skdb_address
The address of an SKDB control block or the symbol for the address.

skmb_address
The address of an SKMB control block or the symbol for the address.

size
The amount of data to display. If the size is greater than the size of the header, the variable portion of the header is displayed if it exists. Must be one to three hexadecimal digits.

HELP
Display IPHDR usage and syntax information instead of the control blocks.

ALL
Display function, operands, and syntax information for the IPHDR command. ALL is the default.

FUNCTION
Display only function information.

OPERANDS
Display only operand information.

SYNTAX
Display only syntax information.

Rule: If you specify multiple keywords from the set {ALL, FUNCTION, OPERANDS, SYNTAX}, all of the keywords that you specify are used.

Sample output of the ICMPHDR command

Following is sample output of the ICMPHDR command.

```
ICMPHDR 9D77428 256

  ICMPv6
   Type/Code        : ECHO Request     CheckSum: 4F51 0000
   Id               : 0028             Seq: 0
   Time             : 2002/05/23 18:43:00.332756
   Echo Data        : -8
 000000 3CED3834 000513D4 08090A0B 0C0D0E0F |<.84............|
 000010 10111213 14151617 18191A1B 1C1D1E1F |................|
 000020 20212223 24252627 28292A2B 2C2D2E2F | !"#$%&'()*+,-./|
 000030 30313233 34353637 38393A3B 3C3D3E3F |0123456789:;<=>?|
 000040 40414243 44454647 48494A4B 4C4D4E4F |@ABCDEFGHIJKLMNO|
 000050 50515253 54555657 58595A5B 5C5D5E5F |PQRSTUVWXYZ.\.^_|
 000060 60616263 64656667 68696A6B 6C6D6E6F |`abcdefghijklmno|
 000070 70717273 74757677 78797A7B 7C7D7E7F |pqrstuvwxyz{|}~.|
 000080 80818283 84858687 88898A8B 8C8D8E8F |................|
 000090 90919293 94959697 98999A9B 9C9D9E9F |................|
 0000A0 A0A1A2A3 A4A5A6A7 A8A9AAAB ACADAEAF |................|
 0000B0 B0B1B2B3 B4B5B6B7 B8B9BABB BCBDBEBF |................|
 0000C0 C0C1C2C3 C4C5C6C7 C8C9CACB CCCDCECF |................|
 0000D0 D0D1D2D3 D4D5D6D7 D8D9DADB DCDDDEDF |................|
 0000E0 E0E1E2E3 E4E5E6E7 E8E9EAEB ECEDEEEF |................|
 0000F0 F0F1F2F3 F4F5F6F7 F8F9FAFB FCFDFEFF |................|

 Protocol Header    : 8
```

```
000000 80004F51 00280000

Data               : 590    Data Length: 0
000000 3CED3834 000513D4 08090A0B 0C0D0E0F |.......M........|
000010 10111213 14151617 18191A1B 1C1D1E1F |................|
000020 20212223 24252627 28292A2B 2C2D2E2F |................|
```

IPHDR command

Use the IPHDR command to display the IP header fields.

Syntax

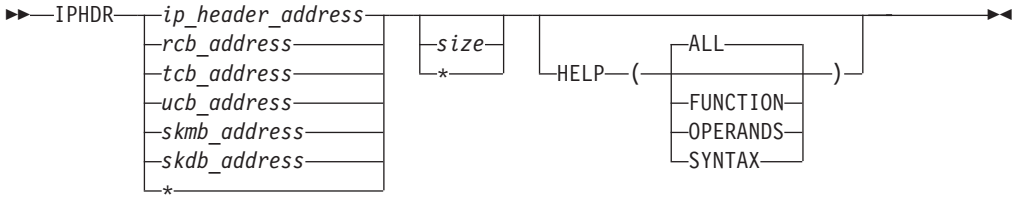

Parameters

***** To omit this positional parameter when using the HELP keyword.

ip_header_address
The address of an IP header or the symbol for the address.

rcb_address
The address of a raw control block or the symbol for the address.

tcb_address
The address of a TCP/IP TCB control block or the symbol for the address.

ucb_address
The address of a UDP control block or the symbol for the address.

skmb_address
The address of an SKMB control block or the symbol for the address.

skdb_address
The address of an SKDB control block or the symbol for the address.

size
The amount of data to display. If the size is greater than the size of the header, additional protocol headers (if any) are displayed. Must be one to three hexadecimal digits.

HELP
Display IPHDR usage and syntax information instead of the control blocks.

ALL
Display function, operands, and syntax information for the IPHDR command. ALL is the default.

FUNCTION
Display only function information.

OPERANDS
Display only operand information.

SYNTAX
Display only syntax information.

Rule: If you specify multiple keywords from the set {ALL, FUNCTION, OPERANDS, SYNTAX}, all of the keywords that you specify are used.

Sample output of the IPHDR command

The following is a sample output of the IPHDR command:

```
IPHDR 09D77400 300

  IP Header: 09D77400
   IpHeader: Version : 6               Header Length: 40
    Class:           : 00              Flow: 000000
    Payload Length   : 264
    Hops             : 255             Protocol: ICMPv6
    Source           : ::1
    Destination      : ::1

   ICMPv6
    Type/Code        : ECHO Request    CheckSum: 4F51 0000
    Id               : 0028            Seq: 0
    Time         ed    : 2002/05/23 18:43:00.332756
    Echo Data        : 256
000000 3CED3834 000513D4 08090A0B 0C0D0E0F |<.84............|
000010 10111213 14151617 18191A1B 1C1D1E1F |................|
000020 20212223 24252627 28292A2B 2C2D2E2F | !"#$%&'()*+,-./|
000030 30313233 34353637 38393A3B 3C3D3E3F |0123456789:;<=>?|
000040 40414243 44454647 48494A4B 4C4D4E4F |@ABCDEFGHIJKLMNO|
000050 50515253 54555657 58595A5B 5C5D5E5F |PQRSTUVWXYZ.\.^_|
000060 60616263 64656667 68696A6B 6C6D6E6F |`abcdefghijklmno|
000070 70717273 74757677 78797A7B 7C7D7E7F |pqrstuvwxyz{|}~.|
000080 80818283 84858687 88898A8B 8C8D8E8F |................|
000090 90919293 94959697 98999A9B 9C9D9E9F |................|
0000A0 A0A1A2A3 A4A5A6A7 A8A9AAAB ACADAEAF |................|
0000B0 B0B1B2B3 B4B5B6B7 B8B9BABB BCBDBEBF |................|
0000C0 C0C1C2C3 C4C5C6C7 C8C9CACB CCCDCECF |................|
0000D0 D0D1D2D3 D4D5D6D7 D8D9DADB DCDDDEDF |................|
0000E0 E0E1E2E3 E4E5E6E7 E8E9EAEB ECEDEEEF |................|
0000F0 F0F1F2F3 F4F5F6F7 F8F9FAFB FCFDFEFF |................|

  IP Header          : 40
000000 60000000 01083AFF 00000000 00000000  00000000 00000001 00000000 0000000
000020 00000000 00000001

  Protocol Header    : 8
000000 80004F51 00280000

  Data               : 720    Data Length: 256
000000 3CED3834 000513D4 08090A0B 0C0D0E0F |.......M........ <.84............|
000010 10111213 14151617 18191A1B 1C1D1E1F |................ ................|
000020 20212223 24252627 28292A2B 2C2D2E2F |................  !"#$%&'()*+,-./|
000030 30313233 34353637 38393A3B 3C3D3E3F |................ 0123456789:;<=>?|
000040 40414243 44454647 48494A4B 4C4D4E4F | .........¢.<(+| @ABCDEFGHIJKLMNO|
000050 50515253 54555657 58595A5B 5C5D5E5F |&.........!$*);^ PQRSTUVWXYZ.\.^_|
000060 60616263 64656667 68696A6B 6C6D6E6F |-/.........,%_>? `abcdefghijklmno|
000070 70717273 74757677 78797A7B 7C7D7E7F |.........`:#@'=" pqrstuvwxyz{|}~.|
000080 80818283 84858687 88898A8B 8C8D8E8F |.abcdefghi...... ................|
000090 90919293 94959697 98999A9B 9C9D9E9F |.jklmnopqr...... ................|
0000A0 A0A1A2A3 A4A5A6A7 A8A9AAAB ACADAEAF |..stuvwxyz...... ................|
0000B0 B0B1B2B3 B4B5B6B7 B8B9BABB BCBDBEBF |................ ................|
0000C0 C0C1C2C3 C4C5C6C7 C8C9CACB CCCDCECF |.ABCDEFGHI...... ................|
0000D0 D0D1D2D3 D4D5D6D7 D8D9DADB DCDDDEDF |.JKLMNOPQR...... ................|
0000E0 E0E1E2E3 E4E5E6E7 E8E9EAEB ECEDEEEF |..STUVWXYZ...... ................|
0000F0 F0F1F2F3 F4F5F6F7 F8F9FAFB FCFDFEFF |0123456789...... ................|
```

RESOLVER command

Use the RESOLVER command to format and summarize resolver control blocks and cache information.

Syntax

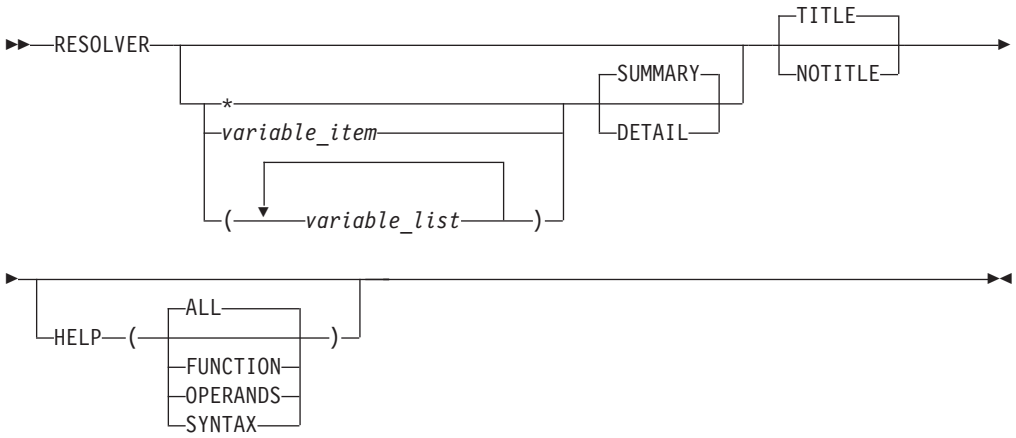

Parameters

If no parameters are specified, information about the Resolver is summarized.

***** An asterisk is used as a placeholder if no variable parameters are specified.

variable_item
Any one of the following variable parameters.

variable_list
You can repeat from 1 - 32 of the following variable parameters, each separated by a blank space, within parentheses:

Variable parameters are:

jobname
Displays only the Resolver control blocks for this job name. The job name can be a TCP/IP application name or a stack name. Must be from 1-8 characters.

ASCB_address
Displays the Resolver control blocks with this address space control block (ASCB) address. An IPCS symbol name can be specified for the address. The address is specified as 1-8 hexadecimal digits. If an address begins with digit A–F, prefix the address with a zero to avoid the address being interpreted as a symbol name or as a character string.

ASID_number
Displays the Resolver control blocks with this Address Space Identifier (ASID). The ASID is a hexadecimal number containing one to four digits.

Restriction: To display the resolver cache information, the resolver address space must be included in the dump and be specified as one of the address spaces to be displayed.

In addition to the variable parameters described above, you can describe the following keyword parameters:

HELP
Display RESOLVER usage and syntax information instead of the control blocks.

ALL
Displays help about the function, operands, and syntax information for the RESOLVER command. ALL is the default.

FUNCTION
Display only function help information.

OPERANDS
Display only operands help information.

SYNTAX
Display only syntax help information.

SUMMARY
Displays the addresses of the control blocks and other data in tables. SUMMARY is the default.

DETAIL
In addition to the SUMMARY display, DETAIL also shows the contents of the control blocks.

TITLE
The title contains information about the dump and about the RESOLVER command. The title information is displayed as the default. The title contains the following information:

- RESOLVER command input parameters.
- Dump data set name.
- Dump title.

NOTITLE
Suppress the title lines. Use this when you are processing lots of commands on the same dump and do not need to see the title information repeated.

Rules:

- If you specify multiple keywords from the set {TITLE,NOTITLE}, only the last one is used.
- If you specify multiple keywords from the set {SUMMARY, DETAIL}, only the last one is used.
- If you specify multiple keywords from the set {ALL, FUNCTION, OPERANDS, SYNTAX}, all of the keywords that you specify are used.

Sample output of the RESOLVER command

The following is sample output of the RESOLVER command with only DETAIL specified:

```
RESOLVER * DETAIL
 Dataset: D74L.AA112546.DUMP3
 Title:   IPCS RESOLVER * DETAIL EXAMPLE

 ================================================================================

 Resolver Analysis

  RCRT: 019F9060
  +0000  RCRTENT#. 0000001F  RCRTRCVT. 9571E000

         -- Array elements --
  +0008  RCRTENTS. 019F9130  RCRTENTS. 019F92BC  RCRTENTS. 019F9448  RCRTENTS. 019F95D4  RCRTENTS. 019F9760  RCRTENTS. 019F98EC
```

```
+0020  RCRTENTS. 019F9A78  RCRTENTS. 019F9C04  RCRTENTS. 019F9D90  RCRTENTS. 019F9F1C  RCRTENTS. 019FA0A8  RCRTENTS. 019FA234
+0038  RCRTENTS. 019FA3C0  RCRTENTS. 019FA4BC  RCRTENTS. 019FA5B8  RCRTENTS. 019FA744  RCRTENTS. 019FA8D0  RCRTENTS. 019FAA5C
+0050  RCRTENTS. 019FABE8  RCRTENTS. 019FAD74  RCRTENTS. 019FAF00  RCRTENTS. 019FAFFC  RCRTENTS. 94BB7180  RCRTENTS. 01083010
+0068  RCRTENTS. 01083010  RCRTENTS. 01083010  RCRTENTS. 01083010  RCRTENTS. 01083010  RCRTENTS. 01083010  RCRTENTS. 019FB0F8
       -- End of array   --

RCVT: 9571E000
+0000  RCVTID......... RCVT
+0004  RCVTASCB....... 00F8C400
+0008  RCVTETKN....... 7FF890D0
+000C  RCVTLX......... 00002700
+0010  RCVTREFR....... 00000007
+0014  RCVTTCA........ 7FF5A868
+0018  RCVTTOPT....... F9000000
+001C  RCVTRSMT....... 15684000
+0020  RCVTDEFA....... 00000001
+0024  RCVTDEFI....... 00000002
+0028  RCVTGBLA....... 00000003
+002C  RCVTGBLI....... 00000004

       -- Array elements --
      RCVTCFG........
+0030  ................................................................................................................................
+00AC  ................................................................................................................................
+0128  ........
      RCVTCFG........
+0130  USER1.OUR.DEFAULT.TCPDATA..................................................................................................
+01AC  ................................................................................................................................
+0228  ........
      RCVTCFG........
+0230  /etc/TRDP3461_resolv.conf.....................................................................................................
+02AC  ................................................................................................................................
+0328  ........
      RCVTCFG........
+0330  ................................................................................................................................
+03AC  ................................................................................................................................
+0428  ........
      RCVTCFG........
+0430  SYS1.TCPPARMS(IPNODES).....................................................................................................
+04AC  ................................................................................................................................
+0528  ........
      RCVTCFG........
+0530  ................................................................................................................................
+05AC  ................................................................................................................................
+0628  ........
      RCVTCFG........
+0630  ................................................................................................................................
+06AC  ................................................................................................................................
+0728  ........
      RCVTCFG........
+0730  /tmp/IPv4Addresses_IPv6Addresses_@mappingsforallSysplexHostnames.....................................................
+07AC  ................................................................................................................................
+0828  ........
      RCVTCFG........
+0830  SYS1.TCPPARMS(SETUPRES).....................................................................................................
+08AC  ................................................................................................................................
+0928  ........
       -- End of array   --
+0930  RCVTDEFD....... 00000000
+0934  RCVTGBLD....... 00000000
+0938  RCVTDEFLHOSTD.. 00000000
+093C  RCVTGBLLHOSTD.. 00000000
+0940  RCVTCART....... 00000000  00000000
+0948  RCVTCNID....... 00000001
+094C  RCVTINID....... 00000000
+0950  RCVTABND....... 00000000
+0954  RCVTJOBN....... RESOLVER
+095C  RCVTSTIM....... C60B6057  BD64BCA6
+0964  RCVTPTIM....... 00000000  00000000
+096C  RCVTCSELOCK.... 00
+096E  RCVTUNRESP..... 0019
+0970  RCVTDEFLHOSTA.. 00000005
+0974  RCVTDEFLHOSTI.. 00000006
+0978  RCVTGBLLHOSTA.. 00000007
+097C  RCVTGBLLHOSTI.. 00000008
+0980  RCVTFLAGSCS.... C0000001
+0984  RCVTSYM@....... 860DC000
+0988  RCVTRIMT....... 98000000
+098C  RCVTCSED....... 00000000

       -- Array elements --
      RCVTNSCTABLE...
+0990  00000000  00000000  00000000  0A000601  00000000  00000000  00000000  00000000  000000D8  00000000  00000000  00000000
+09C0  00000000  00000000  00000004  00000000  00000008  00000000  00000002  00000000  00000005  00000000  00000002  00000000
+09F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0A20  00000000  00000000  00000000  00000000
      RCVTNSCTABLE...
+0A30  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0A60  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0A90  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
```

```
+0AC0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0AD0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0B00  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0B30  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0B60  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0B70  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0BA0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0BD0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0C00  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0C10  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0C40  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0C70  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0CA0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0CB0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0CE0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0D10  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0D40  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0D50  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0D80  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0DB0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0DE0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0DF0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E20  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E50  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E80  00000000  00000000  00000000  00000000

       RCVTNSCTABLE...
+0E90  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0EC0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0EF0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0F20  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0F30  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0F60  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0F90  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0FC0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+0FD0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1030  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1060  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1070  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+10A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+10D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1100  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1110  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1140  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1170  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+11A0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+11B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+11E0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1210  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1240  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1250  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1280  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+12B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+12E0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+12F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1320  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1350  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1380  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1390  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+13C0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+13F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1420  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1430  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1460  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1490  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+14C0  00000000  00000000  00000000  00000000

       RCVTNSCTABLE...
+14D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1500  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1530  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1560  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1570  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+15A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
```

```
+15D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1600  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1610  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1640  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1670  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+16A0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+16B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+16E0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1710  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1740  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1750  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1780  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+17B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+17E0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+17F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1820  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1850  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1880  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1890  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+18C0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+18F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1920  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1930  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1960  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1990  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+19C0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+19D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1A00  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1A30  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1A60  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1A70  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1AA0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1AD0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1B00  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1B10  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1B40  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1B70  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1BA0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1BB0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1BE0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1C10  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1C40  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1C50  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1C80  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1CB0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1CE0  00000000  00000000  00000000  00000000
       RCVTNSCTABLE...
+1CF0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1D20  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1D50  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+1D80  00000000  00000000  00000000  00000000
       -- End of array   --

+1D90  RCVTCACHELIMIT.... 00000000  00000000
+1D98  RCVTCACHECONST.... 00000000  00000000
+1DA0  RCVTCACHECRIT..... 00000000  00000000
+1DA8  RCVTCACHEEXHST.... 00000000  00000000
+1DB0  RCVTCACHESIZE..... 00000000  00000000
+1DB8  RCVTHTINSTID...... 00000000  00000000
+1DC0  RCVTHTABCPID...... 00000000  00000000
+1DC8  RCVTHTELCPID...... 00000000  00000000
+1DD0  RCVTNSRVCPID...... 00000000  00000000
+1DD8  RCVTLKUPCPID...... 00000000  00000000
+1DE0  RCVTNSRVTOKEN..... 00000000  00000000  00000000  00000000
+1DF0  RCVTMAXTTL........ 00000000
+1DF4  RCVTMAXNEGTTL..... 00000000
+1DF8  RCVTCACHEREFS..... 00000000  00000000
+1E00  RCVTCACHEHITS..... 00000000  00000000
+1E08  RCVTLSTOKEN....... 7F5E8D00  000000F2
+1E10  RCVTCACHEHWM...... 00000000  00000000
+1E18  RCVTCACHENEGSZ.... 00000000  00000000
+1E20  RCVTCACHENEGCP.... 00000000  00000000
+1E28  RCVTXMQ........... 00000000  00000000
+1E30  RCVTGTD@.......... 7F553FF0
+1E34  RCVTOPTIONS....... 40A00000
+1E38  RCVTGTDCKSM....... 8301F250  ACAF13C2
+1E40  RCVTNSPT@......... 7E51D9E0
+1E44  RCVTUNRESPDESIRED. 0000
```

```
  +1E48  RCVTMONMSG........ 00000000  00000000
  +1E48  RCVTNSROUTD....... 00000000
  +1E4C  RCVTUPSOUTD....... 00000000
  +1E50  RCVTOMVSEXITECB... 807FF1E8
  +1E54 RCVTMODGBLDECB.... 00000000
  +1E58 RCVT_HC_RETRANS... 0000
Resolver Services Module Table
 Module Table entries by Entry Point Address

 Module   Epa      Date       Time     PTF      Lmod     Asid
 EZBRESMT 15684000 2010/05/25 19:56:00 HIP61D0  EZBRESRV 0001
 EZBRESRV 15684498 2010/05/27 16:28:00 HIP61D0  EZBRESRV 0001
 EZBREARR 15691678 2010/05/27 16:24:00 HIP61D0  EZBRESRV 0001
 EZBRECSR 156929C0 2010/05/27 16:24:00 HIP61D0  EZBRESRV 0001

        ...  edited out part of output  ...

 EZBRETRM 15718238 2010/05/27 16:28:00 HIP61D0  EZBRESRV 0001
 EZBRETRR 15718438 2010/05/27 16:28:00 HIP61D0  EZBRESRV 0001

 Resolver Address Space is RESOLVER (ASID 000F)

 Resolver setup file: SYS1.TCPPARMS(SETUPRES)
 Resolver setup file warning messages: No
 Global TCPIP.DATA file: /etc/TRDP3461_resolv.conf
 Default TCPIP.DATA file: None
 Global IPNODES file: None
 Default IPNODES file: SYS1.TCPPARMS(IPNODES)
 Common Search: Yes
 Unresponsive Threshold: 25
 Autoquiesce  : Yes
 Caching      : Yes
 Cache Limit  : 209715200
 Cache Size   : 80352
 Cache Status : Unconstrained
 Cache Refs   : 28
 Cache Hits   : 0
 Cache Max Use: 80352
 Max TTL      : 2147483647
 Max Neg TTL  : 2147483647
 Cache Reorder: YES

 Resolver Initialization Module Table

 Module Table entries by Entry Point Address
 Module   Epa      Date       Time     PTF      Lmod     Asid
 EZBREIMT 18000000 2010/05/25 19:54:00 HIP61D0  EZBREINI 000F
 EZBREINI 180002F8 2010/05/27 16:07:00 HIP61D0  EZBREINI 000F

        ...  edited out part of output  ...

 EZBRETRC 1802B258 2010/05/27 16:28:00 HIP61D0  EZBREINI 000F
 EZBRETRM 1802BCA8 2010/05/27 16:28:00 HIP61D0  EZBREINI 000F

  CTCA: 7FF5A868
  +0000  CTCA_CBID....... CTCA
  +0004  CTCA_CBSIZE..... 0398
  +0006  CTCA_CBVER...... 0001
  +0008  CTCA_CTNAME..... C5E9C2C3  C3E3D9C3  00F8C400  00000000
  +0018  CTCA_CTTOKEN.... C60B6057  7FF5A868  15E5B300  00F8C400
  +0028  CTCA_CURCTBS.... 7E5BCB58

        ...  edited out part of output  ...

  +0384  CTCA_OPTFLAG.... 00000000
  +0388  CTCA_OPTNAME.... ........
  +0390  CTCA_OPTFLAG.... 00000000
          -- End of array   --
Name Server Polling Task

 NSPT: 7E51D9E0
 +0000  NSPTID.............. NSPT
 +0004  NSPTFLAGS.......... C0000000
        NSPTQUERYBUFFER....
 +0008  91370000  00010000  00000000  03455A42  0653414D  504C4503  434F4D00
 +0024  00010001  00000000  00000000  00000000  00000000  00000000  00000000
 +0040  00000000  00000000
        NSPT_STATE.........
 +0048  00000001  00000001  00000281  00000001  10020035  092A6A02  00000000
 +0064  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0080  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +009C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00B8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00D4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +010C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0128  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0144  00000000  00000000  00000000  00000000  00000000  91370000  7F565180
 +0160  00000000  00000000  00000000  00000000  00000000  00000000  A383974B
```

```
+017C  99819385  8987884B  8982944B  83969400  00000000  00000000  00000000
+0198  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+01B4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+01D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+01EC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0208  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0224  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0240  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+025C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0278  00000000  10000000  00000000  00000000  00000000  00000000  00000000
+0294  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+02B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+02CC  00000000  00000000  00000000  00000000  00000000  7E52D498  00000000
+02E8  00000002  7E52D2F0
      NSPT_STATE_EXT.....
+02F0  00000000  7F5649B4  00000000  7F5648B4  D4E5E2D9  00000000  00000000
+030C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0328  00000000  00000000  00000000  00000000  00000000  00000000  E3C3D7C3
+0344  E2000000  00C3E2F3  F9F00000  00000000  00000000  00000000  00000000
+0360  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+037C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0398  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+03B4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+03D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+03EC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0408  00D4E5E2  D9000000  00000000  00000000  00000000  00000000  00000000
+0424  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0440  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+045C  00000000  00000000  21200000  7F564270  00000000  00000000  FFFFFFFF
+0478  00000000  00000000  0001FFFF  000003E8  00000000  00000000  00000000
+0494  00000000
      NSPT_STATE_IPV6EXT.
+0498  00000002  1C130035  00000000  11110000  00000000  00000000  00006666
+04B4  00000000  10020035  0908080C  00000000  00000000  00000000  00000000
+04D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+04EC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0508  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0524  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0540  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+055C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0578  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0594  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+05B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+05CC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+05E8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0604  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0620  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+063C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0658  00000000
       -- Array elements --
      NSPTNSTABLE........
+0660  00000000  00000000  7E52D760  0000001C  00000000  00000000  00000000
+067C  00000000  00000000  00000000  00000000  00000000  000000F3  00010000
+0698  00000000  00000000  00000000  0000044F  00000000  00000000  18034FD0
+06B4  00000001  7E4FC3E0  00000000  00000000  00000000  00000000  00000000
+06D0  00000000  00000100  00701700  0000000B  00000000  00000000  00000000
+06EC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0708  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0724  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0740  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+075C  00000000  7E4FC98C  0000001C  7E52D788  00000001  00000000  00000000
+0778  00000000  00000001  7E4FC3E0  00000000  7E52D7B4  0000000C  00000000
+0794  00000000  1C130035  00000000  11110000  00000000  00000000  00006666
+07B0  00000000  00000000  00000000  00000000  EC000000  00000001  0000000A
+07CC  00000000  00000000  00000000  000C6BA9  E0AF5E6E  91418000  00000000
+07E8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0804  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0820  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+083C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0858  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0874  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0890  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+08AC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
      NSPTNSTABLE........
+08C8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+08E4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0900  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+091C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0938  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0954  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0970  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+098C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+09A8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+09C4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+09E0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+09FC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0A18  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0A34  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0A50  00000000  00000000  00000000  00000000  00000000  00000000  00000000
```

```
+0A6C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0A88  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0AA4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0AC0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0ADC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0AF8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0B14  00000000  00000000  00000000  00000000  00000000  00000000  00000000
       NSPTNSTABLE........
+0B30  00000001  00000000  7E52DC30  0000001C  00000000  00000000  00000000
+0B4C  00000000  00000000  00000000  00000000  00000000  000000F3  00010000
+0B68  00000000  00000000  00000000  0000044F  00000000  00000000  18034FD0
+0B84  00000003  7E4FC3E0  00000000  00000000  00000000  00000000  00000000
+0BA0  00000000  00000100  00701A00  00000001  00000000  00000000  00000000
+0BBC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0BD8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0BF4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0C10  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0C2C  00000000  7E4FC9C4  00000010  7E52DC58  00000001  00000000  00000000
+0C48  00000000  00000003  7E4FC3E0  00000000  7E52DC84  0000000C  00000000
+0C64  00000000  10020035  0908080C  00000000  00000000  00000000  00000000
+0C80  00000000  00000000  00000000  00000000  EC000000  00000001  0000000A
+0C9C  00000001  00000000  00000000  000C6BA9  E0AF6337  91418000  00000000
+0CB8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0CD4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0CF0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0D0C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
.
.

+5050  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+506C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5088  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+50A4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+50C0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+50DC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
       NSPTNSTABLE........
+50F8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5114  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5130  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+514C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5168  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5184  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+51A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+51BC  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+51D8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+51F4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5210  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+522C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5248  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5264  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5280  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+529C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+52B8  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+52D4  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+52F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+530C  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5328  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+5344  00000000  00000000  00000000  00000000  00000000  00000000  00000000
        -- End of array   --


  Resolver Cache
  Nameserver Cache
  Hash Table Header   : 00000048_00001000
  Instance            : 1
  Active entries      : 3
  Hash buckets        : 251

Index Element@          IP Address
                        IPv4 Fwlk Table@  IPv6 Fwlk Table@  Rvlk Table@
   63 00000048_00108400 127.0.0.1
                        00000000_00000000 00000000_00000000 00000000_00000000
   65 00000048_00108200 10.11.25.46
                        00000000_00000000 00000000_00000000 00000000_00000000
  152 00000048_00108000 10.0.6.1
                        00000048_00005000 00000048_00003000 00000048_00007000

IPv4 Forward Lookup Cache for Nameserver 10.0.6.1
  Hash Table Header   : 00000048_00005000
  Instance            : 3
  Active entries      : 11
  Hash buckets        : 251

Index Element@          Hostname                        Rcode    IP Address
    0 00000048_00109C00 BOGUSNAME.IBM.COX.IBM.COM       NXDOMAIN 0.0.0.0
   26 00000048_00109800 BOGUSNAME.IBM.COX.TCP.RALEIGH.IB NXDOMAIN 0.0.0.0
                        M.COM
   35 00000048_00109A00 BOGUSNAME.IBM.COX.RALEIGH.IBM.CO NXDOMAIN 0.0.0.0
                        M
   39 00000048_00109600 BOGUSNAME.IBM.COX               NXDOMAIN 0.0.0.0
```

```
    69 00000048_00108A00 WWW.IETF.ORG                     NOERROR  64.170.98.32
    86 00000048_0010AC00 BOGUSNAME2                       NXDOMAIN 0.0.0.0
   108 00000048_00108800 WWW.YAHOO.COM                    NOERROR  69.147.125.65
   124 00000048_0010A400 BOGUSNAME2.RALEIGH.IBM.COM       NXDOMAIN 0.0.0.0
   134 00000048_00109000 WWW.IBM.COM                      NOERROR  129.42.60.216
   172 00000048_0010A800 BOGUSNAME2.IBM.COM               NXDOMAIN 0.0.0.0
   173 00000048_0010A000 BOGUSNAME2.TCP.RALEIGH.IBM.COM   NXDOMAIN 0.0.0.0

 IPv6 Forward Lookup Cache for Nameserver 10.0.6.1
   Hash Table Header   : 00000048_00003000
   Instance            : 2
   Active entries      : 7
   Hash buckets        : 251

 Index Element@          Hostname                       Rcode    IP Address
    69 00000048_00109200 WWW.IETF.ORG                   NOERROR  2001:1890:1112:1::20
    86 00000048_0010AA00 BOGUSNAME2                     NXDOMAIN ::
   108 00000048_00108600 WWW.YAHOO.COM                  NXDOMAIN ::
   124 00000048_0010A200 BOGUSNAME2.RALEIGH.IBM.COM     NXDOMAIN ::
   134 00000048_00108C00 WWW.IBM.COM                    NXDOMAIN ::
   172 00000048_0010A600 BOGUSNAME2.IBM.COM             NXDOMAIN ::
   173 00000048_00109E00 BOGUSNAME2.TCP.RALEIGH.IBM.COM NXDOMAIN ::
 Reverse Lookup Cache for Nameserver 10.0.6.1
   Hash Table Header   : 00000048_00007000
   Instance            : 4
   Active entries      : 2
   Hash buckets        : 251

 Index Element@          IP Address                              Rcode    Hostname
   120 00000048_00109400 2001:1890:1112:1::20                    NOERROR  mail.ietf.org
   222 00000048_00108E00 64.170.98.32                            NOERROR  mail.ietf.org

Resolver Process Data for RESOLVER Asid=000F Tcb@=007FF368:
  RPID: 7F5CFFF0
  +0000  RES_LHOSTFILE@........ 00000000
  +0004  RES_LIOWA@............ 7E50EF08
  +0008  RES_PARSE_INDICATOR... FFFFFFFF
          -- Array elements --
  +000C  RES_SHADOW_INDICATOR.. 5E731220
  +0010  RES_SHADOW_INDICATOR.. 00000000
         -- End of array   --
  +0014  RES_CALLER_API........ FFFFFFFF
  +0018  RES_USER_JOB.......... RESOLVER
  +0020  RES_LSERVFILE@........ 00000000
  +0024  RES_IOCINET_TIMESTAMP. 00000000  00000000
  +002C  RES_IOCINETFLAGS...... 0000
  +0030  RPID_IDENTIFIER....... RPID
  +0034  RPID_LENGTH........... 1010
  +0036  RPID_SUBPOOL.......... F9
  +0037  RPID_USERKEY.......... 06
  +0038  RPID_SEQUENCE#........ 00000007
  +003C  RPID_PROCESSID........ 00000002
  +0040  RPIDMGR_TOKEN......... 00000000
  +0044  RPIDMGR_DATA.......... 00000000
  +0048  RPID_RTSK@............ 7F5CF270
  +004C  RPID_RPID@............ 00000000
  +0058  RPID_LATCHHDR@........ 7F5CD000
  +005C  RPID_LATCHSET@........ 7F5CEFF0
         RPID_STATE............
  +0060  00000001  00000002  000002C1  00000003  10020035  0A000601  00000000  00000000  10020035  0A0B192E  00000000  00000000
  +0090  10020035  7F000001  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +00C0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +00F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +0120  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +0150  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  29A20000  7F5D0180  7F5D086C  7F5D096C
  +0180  00000000  00000000  00000000  00000000  A383974B  99819385  8987884B  8982D44B  8396D400  00000000  00000000  00000000
  +01B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +01E0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +0210  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +0240  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +0270  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  24000000  CC925163  FFFF0000
  +02A0  C6851063  00FF0000  C6000000  FF000000  C6851063  C0000000  00000000  00000000  00000000  00000000  00000000  00000000
  +02D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  7F5D04A0  00000000
  +0300  00000002  7F5D02F8
         RPID_STATE_EXT........
  +0308  00000000  7F5CF9B4  00000000  7F5CF8B4  94A5E2F0  F0F00000  00000000  00000000  00000000  00000000  00000000  00000000
  +0338  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  E3C3D7C3  E2000000  00E4E2C5  D9F30000
  +0368  00000000  00000000  00000000  00000000  00000000  00000000  00000009  7F5D03C0  7F5D03C9  7F5D03D2  7F5D03DB  7F5D03E4
  +0398  7F5D03ED  7F5D03F6  7F5D03FF  7F5D0408  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +03C8  00000000  00000000  D1C9E2F7  F8D2D100  00D1C9E2  F8F3D2D1  0000E2D1  C9E2D2C1  D5D100C5  E4C3D2C1  D5D1C900  D2E2C3F5
  +03F8  F6F0F100  00C8C1D5  C7C5E4D3  0000E3C3  C8C9D5C5  E2C500C2  C9C7F500  00000000  E2C3C8C9  D5C5E2C5  00D4E5E2  F0F0F000
  +0428  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +0458  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  21A00000  7F5CF270  00000000
  +0488  00000000  FFFFFFFF  00000000  00000000  0001FFFF  000004D2  00000000  00000000  00000000  00000000
         RPID_STATE_IPV6EXT....
  +04B0  00000006  1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000  10020035  0a000601  00000000  00000000
  +04E0  00000000  00000000  00000000  1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000  10020035  0A0B192E
  +0510  00000000  00000000  00000000  00000000  00000000  10020035  7F000001  00000000  00000000  00000000  00000000  00000000
  +0540  1C130035  00000000  00000000  00000000  00000000  00000001  00000000  00000000  00000000  00000000  00000000  00000000
  +0570  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
```

```
+05A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+05D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0600  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0630  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0660  00000000  00000000  00000000  00000000  00000000
      RPID_ASCII2EBCDIC.....
+0674  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+06A4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+06D4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0704  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0734  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0764  00000000  00000000  00000000  00000000

      RPID_EBCDIC2ASCII.....
+0774  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+07A4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+07D4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0804  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0834  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0864  00000000  00000000  00000000  00000000
+0874  RPID_FLAGS............ C000
+0878  RPID_SEARCH_COUNT..... 00000003
       -- Array elements --        RPID_SEARCH_ENTRY.....
+087C  raleigh.ibm.com.....................................................................................................
+08F8  .....................................................................................................
+0974  ........
      RPID_SEARCH_ENTRY.....
+097C  IbM.cOm.....................................................................................................
+09F8  .....................................................................................................
+0A74  ........
      RPID_SEARCH_ENTRY.....
+0A7C  .....................................................................................................
+0AF8  .....................................................................................................
+0B74  ........
      RPID_SEARCH_ENTRY.....
+0B7C  .....................................................................................................
+0BF8  .....................................................................................................
+0C74  ........
      RPID_SEARCH_ENTRY.....
+0C7C  .....................................................................................................
+0CF8  .....................................................................................................
+0D74  ........
       -- End of array   --
+0D7C  RPID_LATCHTOKEN....... 7FF5A400  00000099
      RPID_TCPIPNAME........
+0D84  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0DB4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0DE4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E14  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E44  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E74  00000000  00000000  00000000  00000000
      RPID_XLATENAME........
+0E84  .....................................................................................................
+0F00  .....................................................................................................
+0F7C  ........
       -- Array elements --
+0F84  RPID_BELOW_AIMAP0..... 00000000
+0F88  RPID_BELOW_AIMAP0..... 00000000
+0F8C  RPID_BELOW_AIMAP0..... 00000000
+0F90  RPID_BELOW_AIMAP0..... 00000000
+0F94  RPID_BELOW_AIMAP0..... 00000000
+0F98  RPID_BELOW_AIMAP0..... 00000000
+0F9C  RPID_BELOW_AIMAP0..... 00000000
+0FA0  RPID_BELOW_AIMAP0..... 00000000
+0FA4  RPID_BELOW_AIMAP0..... 00000000
+0FA8  RPID_BELOW_AIMAP0..... 00000000
+0FAC  RPID_BELOW_AIMAP0..... 00000000
+0FB0  RPID_BELOW_AIMAP0..... 00000000
+0FB4  RPID_BELOW_AIMAP0..... 00000000
+0FB8  RPID_BELOW_AIMAP0..... 00000000
+0FBC  RPID_BELOW_AIMAP0..... 00000000
+0FC0  RPID_BELOW_AIMAP0..... 00000000
       -- End of array   --
       -- Array elements --
+0FC4  RPID_ABOVE_AIMAP0..... 00000000
+0FC8  RPID_ABOVE_AIMAP0..... 00000000
+0FCC  RPID_ABOVE_AIMAP0..... 00000000
+0FD0  RPID_ABOVE_AIMAP0..... 00000000
+0FD4  RPID_ABOVE_AIMAP0..... 00000000
+0FD8  RPID_ABOVE_AIMAP0..... 00000000
+0FDC  RPID_ABOVE_AIMAP0..... 00000000
+0FE0  RPID_ABOVE_AIMAP0..... 00000000
+0FE4  RPID_ABOVE_AIMAP0..... 00000000
+0FE8  RPID_ABOVE_AIMAP0..... 00000000
+0FEC  RPID_ABOVE_AIMAP0..... 00000000
+0FF0  RPID_ABOVE_AIMAP0..... 00000000
+0FF4  RPID_ABOVE_AIMAP0..... 00000000
```

```
+0FF8  RPID_ABOVE_AIMAP@..... 00000000
+0FFC  RPID_ABOVE_AIMAP@..... 00000000
+1000  RPID_ABOVE_AIMAP@..... 00000000
        -- End of array   --
 RPID_STATE NSADDR_LIST contains only IPv4 name server addresses
 RPID_STATE_IPV6EXT STAT_NSADDR_LIST contains the full set of name server addresses

RRST: 7F5D0050
+0000  RETRANS.. 00000001  RETRY.... 00000002  OPTIONS.. 000002C1  NSCOUNT.. 00000003
        -- Array elements --
+0010  NSADDR_LIST........           10020035  0A000601  00000000  00000000
+0020  NSADDR_LIST........           10020035  0A0B192E  00000000  00000000
+0030  NSADDR_LIST........           10020035  7F000001  00000000  00000000
+0040  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0050  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0060  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0070  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0080  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0090  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00A0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00B0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00C0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00D0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00E0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00F0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0100  NSADDR_LIST........           00000000  00000000  00000000  00000000
        -- End of array   --
+0110                     ID....... 29A2
        -- Array elements --
+0114  DNSRCH... 7F5D0180  DNSRCH... 7F5D086C  DNSRCH... 7F5D096C  DNSRCH... 00000000  DNSRCH... 00000000  DNSRCH... 00000000
+012C  DNSRCH... 00000000
        -- End of array   --
+0130  DEFDNAME. tcp.raleigh.ibM.coM.....................................................................................
+01A2           .........................................................................................................
+0214           ............................  PFCODE... 00000000
        -- Array elements --
+0238  SORT_LIST..........           CC925163  FFFF0000
+0240  SORT_LIST..........           C6851063  00FF0000
+0248  SORT_LIST..........           C6000000  FF000000
+0250  SORT_LIST..........           C6851063  C0000000
        -- End of array   --
+0298  RES_IPV6_EXTENSION.           7F5D04A0
+02A0  RES_VERSION........           00000002
+02A4  RES_EXTENSION......           7F5D02F8

RRSX: 7F5D02F8
+0004  STAT_EBCDICTOASCII..... 7F5CF9B4
+000C  STAT_ASCIITOEBCDIC..... 7F5CF8B4
+0010  STAT_HOSTNAME.......... mvS000..........................................................
+0050  STAT_SERVICENAME....... TCPCS....
+0059  STAT_C_DSPRFX.......... USER3.....................
+0078  $__IPDBCSNUM........... 00000009
        -- Array elements --
+007C  $__IP_DBCS_LIST........ 7F5D03C0
+0080  $__IP_DBCS_LIST........ 7F5D03C9
+0084  $__IP_DBCS_LIST........ 7F5D03D2
+0088  $__IP_DBCS_LIST........ 7F5D03DB
+008C  $__IP_DBCS_LIST........ 7F5D03E4
+0090  $__IP_DBCS_LIST........ 7F5D03ED
+0094  $__IP_DBCS_LIST........ 7F5D03F6
+0098  $__IP_DBCS_LIST........ 7F5D03FF
+009C  $__IP_DBCS_LIST........ 7F5D0408
        -- End of array   --

        -- Array elements --
+00C8  STAT_C_DBCS............ JIS78KJ..
+00D1  STAT_C_DBCS............ JIS83KJ..
+00DA  STAT_C_DBCS............ SJISKANJ.
+00E3  STAT_C_DBCS............ EUCKANJI.
+00EC  STAT_C_DBCS............ KSC5601..
+00F5  STAT_C_DBCS............ HANGEUL..
+00FE  STAT_C_DBCS............ TCHINESE.
+0107  STAT_C_DBCS............ BIG5.....
+0110  STAT_C_DBCS............ SCHINESE.
        -- End of array   --
+0119  STAT_NODENAME.......... MVS000...
+0122  STAT_OPTIONS........... 0000
+0128  LE_FUNCTION_DESCRIPTOR. 00000000  00000000  00000000  00000000
+0174  TCP_RES_OPTIONS........ 21A00000
+0178  TCP_RES_TASK@.......... 7F5CF270
+017C  TCP_RES_JOBNAME........ ........
+0184  TCP_RES_SOCKET......... FFFFFFFF
+0188  TCP_RES_RESPLEN........ 00000000
+0190  TCP_SYSNAME_SRC........ 0001
+0192  STAT_SORTLIMIT......... FFFF
+0194  STAT_MILLI_TIME........ 000004D2
+0198  TCP_RES_FAMILY......... 0000
+019A  TCP_RES_TYPE........... 0000

RRS6: 7F5D04A0
```

```
 +0000  STAT_NSCOUNT..... 00000006
         -- Array elements --
 +0004  STAT_NSADDR_LIST. 1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000
 +0020  STAT_NSADDR_LIST. 10020035  0A000601  00000000  00000000  00000000  00000000  00000000
 +003C  STAT_NSADDR_LIST. 1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000
 +0058  STAT_NSADDR_LIST. 10020035  0A0B192E  00000000  00000000  00000000  00000000  00000000
 +0074  STAT_NSADDR_LIST. 10020035  7F000001  00000000  00000000  00000000  00000000  00000000
 +0090  STAT_NSADDR_LIST. 1C130035  00000000  00000000  00000000  00000000  00000001  00000000
 +00AC  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00C8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00E4  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0100  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +011C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0138  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0154  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0170  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +018C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01A8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
         -- End of array   --
Resolver Task Data for RESOLVER Asid=000F Tcb@=007FF368:

 RES_TASK: 7F5CF270
 +0000  RES_IDENTIFIER...... RTSK
 +0004  RES_LENGTH.......... 0D80
 +0006  RES_SUBPOOL......... F9
 +0007  RES_USERKEY......... 06
 +0008  RES_SEQUENCE#....... 00000007
 +0010  RESMGR_TOKEN........ 00000000
 +0014  RESMGR_DATA......... 00000000
 +0018  RES_RTSK@........... 7F5CF270
 +001C  RES_RPID@........... 7F5CFFF0
        RES_STATE...........
 +0030  00000001  00000002  000002C1  00000003  10020035  0A000601  00000000  00000000  10020035  0A0B192E  00000000  00000000
 +0060  10020035  7F000001  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0090  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00C0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0120  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  29A20000  7F5CF3D0  7F5CFAB8  7F5CFBB8
 +0150  00000000  00000000  00000000  00000000  A383974B  99819385  8987884B  8982D44B  8396D400  00000000  00000000  00000000
 +0180  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01E0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0210  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0240  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  24000000  CC925163  FFFF0000
 +0270  C6851063  00FF0000  C6000000  FF000000  C6851063  C0000000  00000000  00000000  00000000  00000000  00000000  00000000
 +02A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  7F5CF6F0  00000000
 +02D0  00000002  7F5CF548
        RES_STATE_EXT.......
 +02D8  00000000  7F5CF9B4  00000000  7F5CF8B4  94A5E2F0  F0F00000  00000000  00000000  00000000  00000000  00000000  00000000
 +0308  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  E3C3D7C3  E2000000  00E4E2C5  D9F30000
 +0338  00000000  00000000  00000000  00000000  00000000  00000000  00000009  7F5CF610  7F5CF619  7F5CF622  7F5CF62B  7F5CF634
 +0368  7F5CF63D  7F5CF646  7F5CF64F  7F5CF658  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0398  00000000  00000000  D1C9E2F7  F8D2D100  00D1C9E2  F8F3D2D1  0000E2D1  C9E2D2C1  D5D100C5  E4C3D2C1  D5D1C900  D2E2C3F5
 +03C8  F6F0F100  00C8C1D5  C7C5E4D3  0000E3C3  C8C9D5C5  E2C500C2  C9C7F500  00000000  E2C3C8C9  D5C5E2C5  00D4E5E2  F0F0F000
 +03F8  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0428  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  21A00000  7F5CF270  00000000
 +0458  00000000  FFFFFFFF  00000000  00000000  0001FFFF  000004D2  00000000  00000000  00000000  00000000

        RES_STATE_IPV6EXT...
 +0480  00000006  1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000  10020035  0A000601  00000000  00000000
 +04B0  00000000  00000000  00000000  1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000  10020035  0A0B192E
 +04E0  00000000  00000000  00000000  00000000  00000000  10020035  7F000001  00000000  00000000  00000000  00000000  00000000
 +0510  1C130035  00000000  00000000  00000000  00000000  00000001  00000000  00000000  00000000  00000000  00000000  00000000
 +0540  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0570  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +05A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +05D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0600  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0630  00000000  00000000  00000000  00000000  00000000
        RES_ASCII2EBCDIC....
 +0644  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0674  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +06A4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +06D4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0704  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0734  00000000  00000000  00000000  00000000
        RES_EBCDIC2ASCII....
 +0744  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0774  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +07A4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +07D4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0804  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0834  00000000  00000000  00000000  00000000
 +0844  RES_SEARCH_COUNT.... 00000003
         -- Array elements --
        RES_SEARCH_ENTRY....
 +0848  raleigh.ibm.com.....................................................................................................
 +08C4  .....................................................................................................................
 +0940  ........
        RES_SEARCH_ENTRY....
```

```
+0948  IbM.cOm.....................................................................................................
+09C4  .....................................................................................................
+0A40  ........
       RES_SEARCH_ENTRY....
+0A48  .....................................................................................................
+0AC4  .....................................................................................................
+0B40  ........
       RES_SEARCH_ENTRY....
+0B48  .....................................................................................................
+0BC4  .....................................................................................................
+0C40  ........
       RES_SEARCH_ENTRY....
+0C48  .....................................................................................................
+0CC4  .....................................................................................................
+0D40  ........
        -- End of array   --
+0D48  RES_LTRACE@......... 007D3250
+0D4C  RES_FLAGS........... 1000
+0D50  RES_LHOSTDATA@...... 00000000
+0D54  RES_SOCKET_ID....... FFFFFFFF
+0D58  RES_PROCESS_TCB@.... 007FF368
+0D5C  RES_AIHOSTDATA@..... 00000000
+0D64  RES_IOCTL_SOCKET_ID. FFFFFFFF
+0D68  RES_THOSTFILE@...... 00000000
+0D6C  RES_THOSTENTRY@..... 00000000
+0D70  RES_RIOWA@.......... 00000000

 RES_STATE NSADDR_LIST contains only IPv4 name server addresses
 RES_STATE_IPV6EXT STAT_NSADDR_LIST contains the full set of name server addresses

RRST: 7F5CF2A0
+0000  RETRANS.. 00000001  RETRY.... 00000002  OPTIONS.. 000002C1  NSCOUNT.. 00000003
        -- Array elements --
+0010  NSADDR_LIST........           10020035  0A000601  00000000  00000000
+0020  NSADDR_LIST........           10020035  0A0B192E  00000000  00000000
+0030  NSADDR_LIST........           10020035  7F000001  00000000  00000000
+0040  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0050  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0060  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0070  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0080  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0090  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00A0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00B0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00C0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00D0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00E0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00F0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0100  NSADDR_LIST........           00000000  00000000  00000000  00000000
        -- End of array   --
+0110                     ID....... 29A2
        -- Array elements --
+0114  DNSRCH... 7F5CF3D0  DNSRCH... 7F5CFAB8  DNSRCH... 7F5CFBB8  DNSRCH... 00000000  DNSRCH... 00000000  DNSRCH... 00000000
+012C  DNSRCH... 00000000
        -- End of array   --
+0130  DEFDNAME. tcp.raleigh.ibM.coM.....................................................................................
+01A2           .....................................................................................................
+0214           ............................  PFCODE... 00000000
        -- Array elements --
+0238  SORT_LIST..........           CC925163  FFFF0000
+0240  SORT_LIST..........           C6851063  00FF0000
+0248  SORT_LIST..........           C6000000  FF000000
+0250  SORT_LIST..........           C6851063  C0000000
        -- End of array   --
+0298  RES_IPV6_EXTENSION.           7F5CF6F0
+02A0  RES_VERSION........           00000002
+02A4  RES_EXTENSION......           7F5CF548
RRSX: 7F5CF548
+0004  STAT_EBCDICTOASCII..... 7F5CF9B4
+000C  STAT_ASCIITOEBCDIC..... 7F5CF8B4
+0010  STAT_HOSTNAME.......... mvS000.....................................................
+0050  STAT_SERVICENAME....... TCPCS....
+0059  STAT_C_DSPRFX.......... USER3.....................
+0078  $__IPDBCSNUM........... 00000009
        -- Array elements --
+007C  $__IP_DBCS_LIST........ 7F5CF610
+0080  $__IP_DBCS_LIST........ 7F5CF619
+0084  $__IP_DBCS_LIST........ 7F5CF622
+0088  $__IP_DBCS_LIST........ 7F5CF62B
+008C  $__IP_DBCS_LIST........ 7F5CF634
+0090  $__IP_DBCS_LIST........ 7F5CF63D
+0094  $__IP_DBCS_LIST........ 7F5CF646
+0098  $__IP_DBCS_LIST........ 7F5CF64F
+009C  $__IP_DBCS_LIST........ 7F5CF658
        -- End of array   --
        -- Array elements --
+00C8  STAT_C_DBCS............ JIS78KJ..
+00D1  STAT_C_DBCS............ JIS83KJ..
+00DA  STAT_C_DBCS............ SJISKANJ.
```

```
 +00E3  STAT_C_DBCS............ EUCKANJI.
 +00EC  STAT_C_DBCS............ KSC5601..
 +00F5  STAT_C_DBCS............ HANGEUL..
 +00FE  STAT_C_DBCS............ TCHINESE.
 +0107  STAT_C_DBCS............ BIG5.....
 +0110  STAT_C_DBCS............ SCHINESE.
         -- End of array   --
 +0119  STAT_NODENAME.......... MVS000...
 +0122  STAT_OPTIONS........... 0000
 +0128  LE_FUNCTION_DESCRIPTOR. 00000000  00000000  00000000  00000000
 +0174  TCP_RES_OPTIONS........ 21A00000
 +0178  TCP_RES_TASK@.......... 7F5CF270
 +017C  TCP_RES_JOBNAME........ ........
 +0184  TCP_RES_SOCKET......... FFFFFFFF
 +0188  TCP_RES_RESPLEN........ 00000000
 +0190  TCP_SYSNAME_SRC........ 0001
 +0192  STAT_SORTLIMIT......... FFFF
 +0194  STAT_MILLI_TIME........ 000004D2
 +0198  TCP_RES_FAMILY......... 0000
 +019A  TCP_RES_TYPE........... 0000

 RRS6: 7F5CF6F0
 +0000  STAT_NSCOUNT..... 00000006

         -- Array elements --
 +0004  STAT_NSADDR_LIST. 1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000
 +0020  STAT_NSADDR_LIST. 10020035  0A000601  00000000  00000000  00000000  00000000  00000000
 +003C  STAT_NSADDR_LIST. 1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000
 +0058  STAT_NSADDR_LIST. 10020035  0A0B192E  00000000  00000000  00000000  00000000  00000000
 +0074  STAT_NSADDR_LIST. 10020035  7F000001  00000000  00000000  00000000  00000000  00000000
 +0090  STAT_NSADDR_LIST. 1C130035  00000000  00000000  00000000  00000000  00000001  00000000
 +00AC  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00C8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00E4  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0100  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +011C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0138  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0154  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0170  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +018C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01A8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
         -- End of array   --

Resolver Process Data for USER28 Asid=0022 Tcb@=007E6B70:

 RPID: 7F5C6FF0
 +0000  RES_LHOSTFILE@........ 7F569780
 +0004  RES_LIOWA@............ 7F5C6180
 +0008  RES_PARSE_INDICATOR... FFFFFFFF
         -- Array elements --
 +000C  RES_SHADOW_INDICATOR.. 5E731220
 +0010  RES_SHADOW_INDICATOR.. 80040000
         -- End of array   --
 +0014  RES_CALLER_API........ 00000005
 +0018  RES_USER_JOB.......... USER2
 +0020  RES_LSERVFILE@........ 7F594800
 +0024  RES_IOCINET_TIMESTAMP. 00000000  00000000
 +002C  RES_IOCINETFLAGS...... 0000
 +0030  RPID_IDENTIFIER....... RPID
 +0034  RPID_LENGTH........... 1010
 +0036  RPID_SUBPOOL.......... F9
 +0037  RPID_USERKEY.......... 06
 +0038  RPID_SEQUENCE#........ 00000007
 +003C  RPID_PROCESSID........ 0500001C
 +0040  RPIDMGR_TOKEN......... 00000208
 +0044  RPIDMGR_DATA.......... 7F5C7038
 +0048  RPID_RTSK@............ 7F5C6270
 +004C  RPID_RPID@............ FF5C6FF0
 +0058  RPID_LATCHHDR@........ 7F5C4000
 +005C  RPID_LATCHSET@........ 7F5C5FF0
        RPID_STATE............
 +0060  00000001  00000002  000002C3  00000003  10020035  0A000601  00000000  00000000  10020035  0A0B192E  00000000  00000000
 +0090  10020035  7F000001  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00C0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0120  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0150  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  AAA40000  7F5C7180  7F5C786C  7F5C796C
 +0180  00000000  00000000  00000000  00000000  A383974B  99819385  8987884B  8982D44B  8396D400  00000000  00000000  00000000
 +01B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01E0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0210  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0240  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0270  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  24000000  CC925163  FFFF0000
 +02A0  C6851063  00FF0000  C6000000  FF000000  C6851063  C0000000  00000000  00000000  00000000  00000000  00000000  00000000
 +02D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  7F5C74A0  00000000
 +0300  00000002  7F5C72F8
        RPID_STATE_EXT........
 +0308  00000000  7F5C69B4  00000000  7F5C68B4  94A5E2F0  F0F00000  00000000  00000000  00000000  00000000  00000000  00000000
 +0338  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  E3C3D7C3  E2000000  00E4E2C5  D9F30000
 +0368  00000000  00000000  00000000  00000000  00000000  00000000  00000009  7F5C73C0  7F5C73C9  7F5C73D2  7F5C73DB  7F5C73E4
```

```
+0398  7F5C73ED  7F5C73F6  7F5C73FF  7F5C7408  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+03C8  00000000  00000000  D1C9E2F7  F8D2D100  00D1C9E2  F8F3D2D1  0000E2D1  C9E2D2C1  D5D100C5  E4C3D2C1  D5D1C900  D2E2C3F5
+03F8  F6F0F100  00C8C1D5  C7C5E4D3  0000E3C3  C8C9D5C5  E2C500C2  C9C7F500  00000000  E2C3C8C9  D5C5E2C5  00D4E5E2  F0F0F000
+0428  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0458  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  29A00000  7F5C6270  00000000
+0488  00000000  FFFFFFFF  00000000  00000000  0001FFFF  000004D2  00000000  00000000  00000000  00000000

      RPID_STATE_IPV6EXT....
+04B0  00000006  1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000  10020035  0A000601  00000000  00000000
+04E0  00000000  00000000  00000000  1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000  10020035  0A0B192E
+0510  00000000  00000000  00000000  00000000  00000000  10020035  7F000001  00000000  00000000  00000000  00000000  00000000
+0540  1C130035  00000000  00000000  00000000  00000000  00000001  00000000  00000000  00000000  00000000  00000000  00000000
+0570  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+05A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+05D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0600  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0630  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0660  00000000  00000000  00000000  00000000  00000000
      RPID_ASCII2EBCDIC.....
+0674  00010203  372D2E2F  1605250B  0C0D0E0F  10111213  3C3D3226  18193F27  221D351F  405A7F7B  5B6C507D  4D5D5C4E  6B604B61
+06A4  F0F1F2F3  F4F5F6F7  F8F97A5E  4C7E6E6F  7CC1C2C3  C4C5C6C7  C8C9D1D2  D3D4D5D6  D7D8D9E2  E3E4E5E6  E7E8E9AD  E0BD5F6D
+06D4  79818283  84858687  88899192  93949596  979899A2  A3A4A5A6  A7A8A9C0  4FD0A107  00010203  372D2E2F  1605250B  0C0D0E0F
+0704  10111213  3C3D3226  18193F27  221D351F  405A7F7B  5B6C507D  4D5D5C4E  6B604B61  F0F1F2F3  F4F5F6F7  F8F97A5E  4C7E6E6F
+0734  7CC1C2C3  C4C5C6C7  C8C9D1D2  D3D4D5D6  D7D8D9E2  E3E4E5E6  E7E8E9AD  E0BD5F6D  79818283  84858687  88899192  93949596
+0764  979899A2  A3A4A5A6  A7A8A9C0  4FD0A107
      RPID_EBCDIC2ASCII.....
+0774  00010203  1A091A7F  1A1A1A0B  0C0D0E0F  10111213  1A0A081A  18191A1A  1C1D1E1F  1A1A1C1A  1A0A171B  1A1A1A1A  1A050607
+07A4  1A1A161A  1A1E1A04  1A1A1A1A  14151A1A  20A6E180  EB909FE2  AB8B9B2E  3C282B7C  26A9AA9C  DBA599E3  A89E2124  2A293B5E
+07D4  2D2FDFDC  9ADDDE98  9DACBA2C  255F3E3F  D78894B0  B1B2FCD6  FB603A23  40273D22  F8616263  64656667  686996A4  F3AFAEC5
+0804  8C6A6B6C  6D6E6F70  71729787  CE93F1FE  C87E7374  75767778  797AEFC0  DA5BF2F9  B5B6FDB7  B8B9E6BB  BCBD8DD9  BF5DD8C4
+0834  7B414243  44454647  4849CBCA  BEE8ECED  7D4A4B4C  4D4E4F50  5152A1AD  F5F4A38F  5CE75354  55565758  595AA085  8EE9E4D1
+0864  30313233  34353637  3839B3F7  F0FAA7FF
+0874  RPID_FLAGS............. C000
+0878  RPID_SEARCH_COUNT..... 00000003
       -- Array elements --
      RPID_SEARCH_ENTRY.....
+087C  raleigh.ibm.com.....................................................................................................
+08F8  .....................................................................................................
+0974  ........
      RPID_SEARCH_ENTRY.....
+097C  IbM.cOm.....................................................................................................
+09F8  .....................................................................................................
+0A74  ........
      RPID_SEARCH_ENTRY.....
+0A7C  .....................................................................................................
+0AF8  .....................................................................................................
+0B74  ........
      RPID_SEARCH_ENTRY.....
+0B7C  .....................................................................................................
+0BF8  .....................................................................................................
+0C74  ........
      RPID_SEARCH_ENTRY.....
+0C7C  .....................................................................................................
+0CF8  .....................................................................................................
+0D74  ........
       -- End of array   --
+0D7C  RPID_LATCHTOKEN....... 7F5C3C00  000000E9
      RPID_TCPIPNAME........
+0D84  6185A383  619985A2  9693A54B  83969586  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0DB4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0DE4  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E14  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E44  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0E74  00000000  00000000  00000000  00000000

      RPID_XLATENAME........
+0E84  USER2.STANDARD.TCPXLBIN.
+0F00
+0F7C
       -- Array elements --
+0F84  RPID_BELOW_AIMAP@..... 00000000
+0F88  RPID_BELOW_AIMAP@..... 00000000
+0F8C  RPID_BELOW_AIMAP@..... 00000000
+0F90  RPID_BELOW_AIMAP@..... 00000000
+0F94  RPID_BELOW_AIMAP@..... 00000000
+0F98  RPID_BELOW_AIMAP@..... 00000000
+0F9C  RPID_BELOW_AIMAP@..... 00000000
+0FA0  RPID_BELOW_AIMAP@..... 00000000
+0FA4  RPID_BELOW_AIMAP@..... 00000000
+0FA8  RPID_BELOW_AIMAP@..... 00000000
+0FAC  RPID_BELOW_AIMAP@..... 00000000
+0FB0  RPID_BELOW_AIMAP@..... 00000000
+0FB4  RPID_BELOW_AIMAP@..... 00000000
+0FB8  RPID_BELOW_AIMAP@..... 00000000
+0FBC  RPID_BELOW_AIMAP@..... 00000000
+0FC0  RPID_BELOW_AIMAP@..... 00000000
       -- End of array   --
       -- Array elements --
+0FC4  RPID_ABOVE_AIMAP@..... 00000000
+0FC8  RPID_ABOVE_AIMAP@..... 00000000
+0FCC  RPID_ABOVE_AIMAP@..... 00000000
```

```
+0FD0  RPID_ABOVE_AIMAP0..... 00000000
+0FD4  RPID_ABOVE_AIMAP0..... 00000000
+0FD8  RPID_ABOVE_AIMAP0..... 00000000
+0FDC  RPID_ABOVE_AIMAP0..... 00000000
+0FE0  RPID_ABOVE_AIMAP0..... 00000000
+0FE4  RPID_ABOVE_AIMAP0..... 18142000
+0FE8  RPID_ABOVE_AIMAP0..... 00000000
+0FEC  RPID_ABOVE_AIMAP0..... 00000000
+0FF0  RPID_ABOVE_AIMAP0..... 00000000
+0FF4  RPID_ABOVE_AIMAP0..... 00000000
+0FF8  RPID_ABOVE_AIMAP0..... 00000000
+0FFC  RPID_ABOVE_AIMAP0..... 00000000
+1000  RPID_ABOVE_AIMAP0..... 00000000
       -- End of array   --

 RPID_STATE NSADDR_LIST contains only IPv4 name server addresses
 RPID_STATE_IPV6EXT STAT_NSADDR_LIST contains the full set of name server addresses

RRST: 7F5C7050
+0000  RETRANS.. 00000001  RETRY.... 00000002  OPTIONS.. 000002C3  NSCOUNT.. 00000003
       -- Array elements --
+0010  NSADDR_LIST........           10020035  0A000601  00000000  00000000
+0020  NSADDR_LIST........           10020035  0A0B192E  00000000  00000000
+0030  NSADDR_LIST........           10020035  7F000001  00000000  00000000
+0040  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0050  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0060  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0070  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0080  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0090  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00A0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00B0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00C0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00D0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00E0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00F0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0100  NSADDR_LIST........           00000000  00000000  00000000  00000000
       -- End of array   --
+0110                     ID....... AAA4
       -- Array elements --
+0114  DNSRCH... 7F5C7180  DNSRCH... 7F5C786C  DNSRCH... 7F5C796C  DNSRCH... 00000000  DNSRCH... 00000000  DNSRCH... 00000000
+012C  DNSRCH... 00000000
       -- End of array   --
+0130  DEFDNAME. tcp.raleigh.ibM.coM.....................................................................................
+01A2           .....................................................................................................
+0214           ............................  PFCODE... 00000000
       -- Array elements --
+0238  SORT_LIST..........           CC925163  FFFF0000
+0240  SORT_LIST..........           C6851063  00FF0000
+0248  SORT_LIST..........           C6000000  FF000000
+0250  SORT_LIST..........           C6851063  C0000000
       -- End of array   --
+0298  RES_IPV6_EXTENSION.           7F5C74A0
+02A0  RES_VERSION........           00000002
+02A4  RES_EXTENSION......           7F5C72F8

RRSX: 7F5C72F8
+0004  STAT_EBCDICTOASCII..... 7F5C69B4
+000C  STAT_ASCIITOEBCDIC..... 7F5C68B4
+0010  STAT_HOSTNAME.......... mvS000.....................................................
+0050  STAT_SERVICENAME....... TCPCS....
+0059  STAT_C_DSPRFX.......... USER3.....................
+0078  $__IPDBCSNUM........... 00000009
       -- Array elements --
+007C  $__IP_DBCS_LIST........ 7F5C73C0
+0080  $__IP_DBCS_LIST........ 7F5C73C9
+0084  $__IP_DBCS_LIST........ 7F5C73D2
+0088  $__IP_DBCS_LIST........ 7F5C73DB
+008C  $__IP_DBCS_LIST........ 7F5C73E4
+0090  $__IP_DBCS_LIST........ 7F5C73ED
+0094  $__IP_DBCS_LIST........ 7F5C73F6
+0098  $__IP_DBCS_LIST........ 7F5C73FF
+009C  $__IP_DBCS_LIST........ 7F5C7408
       -- End of array   --
       -- Array elements --
+00C8  STAT_C_DBCS............ JIS78KJ..
+00D1  STAT_C_DBCS............ JIS83KJ..
+00DA  STAT_C_DBCS............ SJISKANJ.
+00E3  STAT_C_DBCS............ EUCKANJI.
+00EC  STAT_C_DBCS............ KSC5601..
+00F5  STAT_C_DBCS............ HANGEUL..
+00FE  STAT_C_DBCS............ TCHINESE.
+0107  STAT_C_DBCS............ BIG5.....
+0110  STAT_C_DBCS............ SCHINESE.
       -- End of array   --
+0119  STAT_NODENAME.......... MVS000...
+0122  STAT_OPTIONS........... 0000
+0128  LE_FUNCTION_DESCRIPTOR. 00000000  00000000  00000000  00000000
+0174  TCP_RES_OPTIONS........ 29A00000
```

```
 +0178  TCP_RES_TASK@.......... 7F5C6270
 +017C  TCP_RES_JOBNAME........ ........
 +0184  TCP_RES_SOCKET......... FFFFFFFF
 +0188  TCP_RES_RESPLEN........ 00000000
 +0190  TCP_SYSNAME_SRC........ 0001
 +0192  STAT_SORTLIMIT......... FFFF
 +0194  STAT_MILLI_TIME........ 000004D2
 +0198  TCP_RES_FAMILY......... 0000
 +019A  TCP_RES_TYPE........... 0000
 RRS6: 7F5C74A0
 +0000  STAT_NSCOUNT..... 00000006
         -- Array elements --
 +0004  STAT_NSADDR_LIST. 1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000
 +0020  STAT_NSADDR_LIST. 10020035  0A000601  00000000  00000000  00000000  00000000  00000000
 +003C  STAT_NSADDR_LIST. 1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000
 +0058  STAT_NSADDR_LIST. 10020035  0A0B192E  00000000  00000000  00000000  00000000  00000000
 +0074  STAT_NSADDR_LIST. 10020035  7F000001  00000000  00000000  00000000  00000000  00000000
 +0090  STAT_NSADDR_LIST. 1C130035  00000000  00000000  00000000  00000000  00000001  00000000
 +00AC  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00C8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00E4  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0100  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +011C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0138  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0154  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0170  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +018C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01A8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
         -- End of array   --

 LHST: 7F569780
 +0000  LHOST_IDENTIFIER... LHST
 +0004  LHOST_SUBPOOL...... F9
 +0005  LHOST_USERKEY...... 06
 +0008  LHOST_LENGTH....... 187C
 +000A  LHOST_CACHE_LENGTH. 125C
 +000C  LHOST_COUNT........ 00000125
 +0010  LHOST_XFIRST@...... 7F53DBA0
 +0014  LHOST_XLAST@....... 7F53DBA0
 +0018  LHOST_ENTRY@....... 7F53DDA4
 +001C  LHOST_NPAGE@....... 7F567000
 +0020  LHOST_CURPAGE@..... 7F53C000
 +0024  LHOST_FLAGS........ 00040000
 +0028  LHOST_FLAGS2....... 00000000
 Site/Addr Cache

Bucket@  Element@ Data

       ...  edited out part of output  ...

7F569AF8 7F569BD4 Type: SiteInfo6   HostEntry
                  Name: LOOPBACK6
                  Addr: ::1
         7F569BC0 Type: SiteInfo4   HostEntry
                  Name: LOOPBACK
                  Addr: 127.0.0.1
         7F569BAC Type: SiteInfo4   HostEntry
                  Name: LOCALHOST
                  Addr: 127.0.0.1

       ...  edited out part of output  ...
 7F569B3C 7F569BE8 Type: SiteInfo6   HostEntry
                  Name: UNSPECIFIED
                  Addr: 0.0.0.0
 LSRV: 7F594800
 +0000  LSERV_IDENTIFIER... LSRV
 +0004  LSERV_SUBPOOL...... F9
 +0005  LSERV_USERKEY...... 06
 +0006  LSERV_INITFLAGS.... 00
 +0007  LSERV_FLAGS........ 80
 +0008  LSERV_LENGTH....... 07FC
 +000A  LSERV_CACHE_LENGTH. 0690
 +000C  LSERV_COUNT........ 00000074
 +0010  LSERV_XFIRST@...... 7F53F818
 +0014  LSERV_XLAST@....... 7F53F818
 +0018  LSERV_ENTRY@....... 7F53F968
 +001C  LSERV_NPAGE@....... 7F56B000
 +0020  LSERV_CURPAGE@..... 7F568000
 +0024  LSERV_EFIRST@...... 7F59482C
 +0028  LSERV_ELAST@....... 7F53F954

Service Cache

Element@ Port# Protocol ServiceName
7F59482C  9999 tcp      IHVsrvr
7F594840  9999 udp      IHVsrvr
7F594854     7 tcp      echo
7F594868     7 udp      echo
7F59487C     9 tcp      discard
```

```
                      sink
                      null
7F594890     9 udp    discard
                      sink
                      null

       ...  edited out part of output  ...

7F53F940  4035 tcp     rse
7F53F954  2252 tcp     njenet-ssl

AddrInfo Pools
Location Key Raim@     Free# Alloc#
Above      8 18142000     11      0

Resolver Task Data for USER28 Asid=0022 Tcb@=007E6B70:
 RES_TASK: 7F5C6270
 +0000  RES_IDENTIFIER...... RTSK
 +0004  RES_LENGTH.......... 0D80
 +0006  RES_SUBPOOL......... F9
 +0007  RES_USERKEY......... 06
 +0008  RES_SEQUENCE#....... 00000007
 +0010  RESMGR_TOKEN........ 00000000
 +0014  RESMGR_DATA......... 00000000
 +0018  RES_RTSK@........... 7F5C6270
 +001C  RES_RPID@........... 7F5C6FF0
        RES_STATE...........
 +0030  00000001  00000002  000002C3  00000003  10020035  0A000601  00000000  00000000  10020035  0A0B192E  00000000  00000000
 +0060  10020035  7F000001  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0090  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00C0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +00F0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0120  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  AAA40000  7F5C63D0  7F5C6AB8  7F5C6BB8
 +0150  00000000  00000000  00000000  00000000  A383974B  99819385  8987884B  8982D44B  8396D400  00000000  00000000  00000000
 +0180  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01B0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +01E0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0210  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0240  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  24000000  CC925163  FFFF0000
 +0270  C6851063  00FF0000  C6000000  FF000000  C6851063  C0000000  00000000  00000000  00000000  00000000  00000000  00000000
 +02A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  7F5C66F0  00000000
 +02D0  00000002  7F5C6548
        RES_STATE_EXT.......
 +02D8  00000000  7F5C69B4  00000000  7F5C68B4  94A5E2F0  F0F00000  00000000  00000000  00000000  00000000  00000000  00000000
 +0308  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  E3C3D7C3  E2000000  00E4E2C5  D9F30000
 +0338  00000000  00000000  00000000  00000000  00000000  00000000  00000009  7F5C6610  7F5C6619  7F5C6622  7F5C662B  7F5C6634
 +0368  7F5C663D  7F5C6646  7F5C664F  7F5C6658  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0398  00000000  00000000  D1C9E2F7  F8D2D100  00D1C9E2  F8F3D2D1  0000E2D1  C9E2D2C1  D5D100C5  E4C3D2C1  D5D1C900  D2E2C3F5
 +03C8  F6F0F100  00C8C1D5  C7C5E4D3  0000E3C3  C8C9D5C5  E2C500C2  C9C7F500  00000000  E2C3C8C9  D5C5E2C5  00D4E5E2  F0F0F000
 +03F8  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0428  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  29A00000  7F5C6270  00000000
 +0458  00000000  FFFFFFFF  00000000  00000000  0001FFFF  000004D2  00000000  00000000  00000000
        RES_STATE_IPV6EXT...
 +0480  00000006  1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000  10020035  0A000601  00000000  00000000
 +04B0  00000000  00000000  00000000  1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000  10020035  0A0B192E
 +04E0  00000000  00000000  00000000  00000000  00000000  10020035  7F000001  00000000  00000000  00000000  00000000  00000000
 +0510  1C130035  00000000  00000000  00000000  00000000  00000001  00000000  00000000  00000000  00000000  00000000  00000000
 +0540  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0570  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +05A0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +05D0  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0600  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
 +0630  00000000  00000000  00000000  00000000  00000000
        RES_ASCII2EBCDIC....
 +0644  00010203  372D2E2F  1605250B  0C0D0E0F  10111213  3C3D3226  18193F27  221D351F  405A7F7B  5B6C507D  4D5D5C4E  6B604B61
 +0674  F0F1F2F3  F4F5F6F7  F8F97A5E  4C7E6E6F  7CC1C2C3  C4C5C6C7  C8C9D1D2  D3D4D5D6  D7D8D9E2  E3E4E5E6  E7E8E9AD  E0BD5F6D
 +06A4  79818283  84858687  88899192  93949596  979899A2  A3A4A5A6  A7A8A9C0  4FD0A107  00010203  372D2E2F  1605250B  0C0D0E0F
 +06D4  10111213  3C3D3226  18193F27  221D351F  405A7F7B  5B6C507D  4D5D5C4E  6B604B61  F0F1F2F3  F4F5F6F7  F8F97A5E  4C7E6E6F
 +0704  7CC1C2C3  C4C5C6C7  C8C9D1D2  D3D4D5D6  D7D8D9E2  E3E4E5E6  E7E8E9AD  E0BD5F6D  79818283  84858687  88899192  93949596
 +0734  979899A2  A3A4A5A6  A7A8A9C0  4FD0A107
        RES_EBCDIC2ASCII....
 +0744  00010203  1A091A7F  1A1A1A0B  0C0D0E0F  10111213  1A0A081A  18191A1A  1C1D1E1F  1A1A1C1A  1A0A171B  1A1A1A1A  1A050607
 +0774  1A1A161A  1A1E1A04  1A1A1A1A  14151A1A  20A6E180  EB909FE2  AB8B9B2E  3C282B7C  26A9AA9C  DBA599E3  A89E2124  2A293B5E
 +07A4  2D2FDFDC  9ADDDE98  9DACBA2C  255F3E3F  D78894B0  B1B2FCD6  FB603A23  40273D22  F8616263  64656667  686996A4  F3AFAEC5
 +07D4  8C6A6B6C  6D6E6F70  71729787  CE93F1FE  C87E7374  75767778  797AEFC0  DA5BF2F9  B5B6FDB7  B8B9E6BB  BCBD8DD9  BF5DD8C4
 +0804  7B414243  44454647  4849CBCA  BEE8ECED  7D4A4B4C  4D4E4F50  5152A1AD  F5F4A38F  5CE75354  55565758  595AA085  8EE9E4D1
 +0834  30313233  34353637  3839B3F7  F0FAA7FF
 +0844  RES_SEARCH_COUNT.... 00000003
         -- Array elements --
        RES_SEARCH_ENTRY....
 +0848  raleigh.ibm.com..........................................................................................
 +08C4  ..........................................................................................
 +0940  ........
        RES_SEARCH_ENTRY....
 +0948  IbM.cOm..........................................................................................
 +09C4  ..........................................................................................
 +0A40  ........
        RES_SEARCH_ENTRY....
 +0A48  ..........................................................................................
```

```
+0AC4  ..................................................................................................
+0B40  ........
       RES_SEARCH_ENTRY....
+0B48  ..................................................................................................
+0BC4  ..................................................................................................
+0C40  ........
       RES_SEARCH_ENTRY....
+0C48  ..................................................................................................
+0CC4  ..................................................................................................
+0D40  ........
        -- End of array   --
+0D48  RES_LTRACE@......... 007CB3A0
+0D4C  RES_FLAGS........... D000
+0D50  RES_LHOSTDATA@...... 18143000
+0D54  RES_SOCKET_ID....... FFFFFFFF
+0D58  RES_PROCESS_TCB@.... 007E6B70
+0D5C  RES_AIHOSTDATA@..... 7F56C000
+0D64  RES_IOCTL_SOCKET_ID. FFFFFFFF
+0D68  RES_THOSTFILE@...... 00000000
+0D6C  RES_THOSTENTRY@..... 00000000
+0D70  RES_RIOWA@.......... 00000000
 RES_STATE NSADDR_LIST contains only IPv4 name server addresses
 RES_STATE_IPV6EXT STAT_NSADDR_LIST contains the full set of name server addresses

RRST: 7F5C62A0
+0000  RETRANS.. 00000001  RETRY.... 00000002  OPTIONS.. 000002C3  NSCOUNT.. 00000003
        -- Array elements --
+0010  NSADDR_LIST........           10020035  0A000601  00000000  00000000
+0020  NSADDR_LIST........           10020035  0A0B192E  00000000  00000000
+0030  NSADDR_LIST........           10020035  7F000001  00000000  00000000
+0040  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0050  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0060  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0070  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0080  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0090  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00A0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00B0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00C0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00D0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00E0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+00F0  NSADDR_LIST........           00000000  00000000  00000000  00000000
+0100  NSADDR_LIST........           00000000  00000000  00000000  00000000
        -- End of array   --
+0110                     ID....... AAA4
        -- Array elements --
+0114  DNSRCH... 7F5C63D0  DNSRCH... 7F5C6AB8  DNSRCH... 7F5C6BB8  DNSRCH... 00000000  DNSRCH... 00000000  DNSRCH... 00000000
+012C  DNSRCH... 00000000
        -- End of array   --
+0130  DEFDNAME. tcp.raleigh.ibM.coM.....................................................................
+01A2            ..........................................................................................
+0214            ............................  PFCODE... 00000000
        -- Array elements --
+0238  SORT_LIST..........           CC925163  FFFF0000
+0240  SORT_LIST..........           C6851063  00FF0000
+0248  SORT_LIST..........           C6000000  FF000000
+0250  SORT_LIST..........           C6851063  C0000000
        -- End of array   --
+0298  RES_IPV6_EXTENSION.           7F5C66F0
+02A0  RES_VERSION........           00000002
+02A4  RES_EXTENSION......           7F5C6548

RRSX: 7F5C6548
+0004  STAT_EBCDICTOASCII..... 7F5C69B4
+000C  STAT_ASCIITOEBCDIC..... 7F5C68B4
+0010  STAT_HOSTNAME.......... mvS000.............................................
+0050  STAT_SERVICENAME....... TCPCS....
+0059  STAT_C_DSPRFX.......... USER3.......................
+0078  $__IPDBCSNUM........... 00000009
        -- Array elements --
+007C  $__IP_DBCS_LIST........ 7F5C6610
+0080  $__IP_DBCS_LIST........ 7F5C6619
+0084  $__IP_DBCS_LIST........ 7F5C6622
+0088  $__IP_DBCS_LIST........ 7F5C662B
+008C  $__IP_DBCS_LIST........ 7F5C6634
+0090  $__IP_DBCS_LIST........ 7F5C663D
+0094  $__IP_DBCS_LIST........ 7F5C6646
+0098  $__IP_DBCS_LIST........ 7F5C664F
+009C  $__IP_DBCS_LIST........ 7F5C6658
        -- End of array   --
        -- Array elements --
+00C8  STAT_C_DBCS............ JIS78KJ..
+00D1  STAT_C_DBCS............ JIS83KJ..
+00DA  STAT_C_DBCS............ SJISKANJ.
+00E3  STAT_C_DBCS............ EUCKANJI.
+00EC  STAT_C_DBCS............ KSC5601..
+00F5  STAT_C_DBCS............ HANGEUL..
+00FE  STAT_C_DBCS............ TCHINESE.
+0107  STAT_C_DBCS............ BIG5.....
+0110  STAT_C_DBCS............ SCHINESE.
```

```
       -- End of array   --
+0119  STAT_NODENAME......... MVS000...
+0122  STAT_OPTIONS........... 0000
+0128  LE_FUNCTION_DESCRIPTOR. 00000000  00000000  00000000  00000000
+0174  TCP_RES_OPTIONS........ 29A00000
+0178  TCP_RES_TASK@.......... 7F5C6270
+017C  TCP_RES_JOBNAME........ ........
+0184  TCP_RES_SOCKET......... FFFFFFFF
+0188  TCP_RES_RESPLEN........ 00000000
+0190  TCP_SYSNAME_SRC........ 0001
+0192  STAT_SORTLIMIT......... FFFF
+0194  STAT_MILLI_TIME........ 000004D2
+0198  TCP_RES_FAMILY......... 0000
+019A  TCP_RES_TYPE........... 0000
RRS6: 7F5C66F0
+0000  STAT_NSCOUNT..... 00000006
       -- Array elements --
+0004  STAT_NSADDR_LIST. 1C130035  00000000  20020000  00000000  00000000  DEADBEEF  00000000
+0020  STAT_NSADDR_LIST. 10020035  0A000601  00000000  00000000  00000000  00000000  00000000
+003C  STAT_NSADDR_LIST. 1C130035  00000000  20022222  AAAABBBB  CCCCDDDD  EEEEFFFF  00000000
+0058  STAT_NSADDR_LIST. 10020035  0A0B192E  00000000  00000000  00000000  00000000  00000000
+0074  STAT_NSADDR_LIST. 10020035  7F000001  00000000  00000000  00000000  00000000  00000000
+0090  STAT_NSADDR_LIST. 1C130035  00000000  00000000  00000000  00000000  00000001  00000000
+00AC  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+00C8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+00E4  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0100  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+011C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0138  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0154  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0170  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+018C  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
+01A8  STAT_NSADDR_LIST. 00000000  00000000  00000000  00000000  00000000  00000000  00000000
       -- End of array   --

APPLDATA: 18143000
+0000  SERV_ID............ RSRV
+0004  SERV_LENGTH........ 3000
+0006  SERV_SUBPOOL....... FB
+0007  SERV_KEY........... 88
+0008  SERVPRFXSECTION.... 00000000
+000C  SERVCTRLSECTION.... 00000000
       SERVNAMESECTION....
+0010  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
+0040  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000

      ...  edited out part of output  ...

+0544  SERV.............. 00000000  00000000  00000000  00000000
       -- Array elements --
+0558  SERV_ALIASES....... 00000000
+055C  SERV_ALIASES....... 00000000

      ...  edited out part of output  ...

+05E4  SERV_ALIASES....... 00000000
       -- End of array   --

+05F0  HOST_ID............ RHST
+05F4  HOST_LENGTH........ 3000
+05F6  HOST_SUBPOOL....... FB
+05F7  HOST_KEY........... 88
       -- Array elements --
+05F8  H_ADDR_PTRS........ 18143980
+05FC  H_ADDR_PTRS........ 00000000

      ...  edited out part of output  ...

+067C  H_ADDR_PTRS........ 00000000
+0684  H_ADDR_PTRS........ 00000000
       -- End of array   --
       -- Array elements --
+0688  H64_ADDR_PTRS...... 00000000  00000000
+0690  H64_ADDR_PTRS...... 00000000  00000000

      ...  edited out part of output  ...

+0798  H64_ADDR_PTRS...... 00000000  00000000
+07A0  H64_ADDR_PTRS...... 00000000  00000000
       -- End of array   --
+07A8  HOST.............. 181439DE  181437DC  00000002  00000004  181435F8
+07BC  HOST64............ 00000000  181439DE  00000000  18143868  00000000  00000000  00000000  18143688
       -- Array elements --
+07DC  HOST_ALIASES....... 00000000
+07E0  HOST_ALIASES....... 00000000
      ...  edited out part of output  ...

+0860  HOST_ALIASES....... 00000000
+0864  HOST_ALIASES....... 00000000
```

```
        -- End of array   --
        -- Array elements --
  +0868  HOST64_ALIASES..... 00000000  00000000
  +0870  HOST64_ALIASES..... 00000000  00000000

       ...  edited out part of output  ...
   +0970  HOST64_ALIASES..... 00000000  00000000
  +0978  HOST64_ALIASES..... 00000000  00000000
        -- End of array   --
  +0980  HOST_ADDR.......... 40AA6220
        -- Array elements --
  +0984  HOST_ADDRS......... 00000000
  +0988  HOST_ADDRS......... 00000000
        -- End of array   --
        -- Array elements --
  +0990  HOST64_ADDRS....... 00000000  00000000
  +0998  HOST64_ADDRS....... 00000000  00000000
        -- End of array   --
  +09A0  HOSTADDR........... 00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  000000
        HOSTBUF............
  +09C4  F3F24BF9  F84BF1F7  F04BF6F4  4B899560  81848499  4B819997  81009481  89934B89  85A3864B  96998700  002A3CD8  00000000

       ...  edited out part of output  ...

  +2090  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
        HOSTBUFFER.........
  +20B0  AAB18183  00010000  00010000  09626F67  75736E61  6D650369  626D0363  6F780349  624D0363  4F6D0000  010001C0  1E000600
  +20E0  0100002A  30002905  6C656461  32C01E07  646E7374  65616D02  7573C01E  77CEF8C1  00001C20  00000E10  00093A80  0000A8C0
  +2110  00070800  00038400  093A8000  00384000  01518053  00000200  01000037  A7000401  61C05300  00020001  000037A7  00040162
  +2140  C0530000  02000100  0037A700  040163C0  53000002  00010000  37A70004  0164C053  00000200  01000037  A7000401  65C05300
  +2170  00020001  000037A7  00040166  C0530000  02000100  0037A700  040167C0  53000002  00010000  37A70004  0168C053  00000200
  +21A0  01000037  A7000401  69C05300  00020001  000037A7  0004016A  C0530000  02000100  0037A700  04016BC0  53C07F00  01000100
  +21D0  033BFA00  04C62900  04C07F00  1C000100  03449500  10200105  03BA3E00  00000000  00000200  30C08E00  01000100  059F4500
  +2200  04C0E44F  C9C09D00  01000100  05444200  04C02104  0CC0AC00  01000100  05806500  0480080A  5AC0BB00  01000100  069AD100
  +2230  04C0CBE6  0AC0CA00  01000100  0392B700  04C00505  F1C0CA00  1C000100  02DAD800  10200105  00002F00  00000000  00000000
  +2260  0FC0D900  01000100  06D14600  04C07024  04C0E800  01000100  03468500  04803F02  35C0E800  1C000100  02DB5100  10200105
  +2290  00000100  00000000  00803F02  35C0F700  01000100  05877600  04C02494  11000000  00000000  00000000  00000000  00000000

       ...  edited out part of output  ...

  +2C80  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
  +2CB0  HOST_ZONEINTERFACE. 00000000  00000000  00000000  00000000

  RAHD: 7F56C000
  +0000  AIH_ID............... RAIH
  +0004  AIH_LENGTH........... 2CA0
  +0006  AIH_SUBPOOL.......... F9
  +0007  AIH_KEY.............. 06
        AIH_BUF.............
  +0008  86974BA6  87F14B82  4BA88188  96964B83  96940000  00000000  00000000  00000000  00000000  00000000  00000000  00000000

       ...  edited out part of output  ...

  +0E6D  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000  00000000
        AIH_CANONNAME.........
  +0E9D  ..........................................................................................................................
  +0F19  ..........................................................................................................................
  +0F95  ........
  +0F9D  AIH_CANONNAMELEN...... 00
  +0FA0  AIH_IOCSTACKINFOFLAGS. 80
  +0FA4  AIH_IOCSTACKENTRIES... 00000001
        -- Array elements --
  +0FA8  AIH_IOCSTACKNAME...... TCPCS
  +0FB0  AIH_IOCSTACKTDINDEX... 01
  +0FB1  AIH_IOCSTACKFLAGS..... 90
  +0FB2  AIH_IOCSTACKNAME...... ........
  +0FBA  AIH_IOCSTACKTDINDEX... 00
  +0FBB  AIH_IOCSTACKFLAGS..... 00

       ...  edited out part of output  ...

  +10DE  AIH_IOCSTACKNAME...... ........
  +10E6  AIH_IOCSTACKTDINDEX... 00
  +10E7  AIH_IOCSTACKFLAGS..... 00
        -- End of array   --
  +11A8  AIH_SICOUNT........... 00
        -- Array elements --
  +11AC  AIH_SISOCKTYPE........ 00000001
  +11B0  AIH_SIPROTOCOL........ 00000006
  +11B4  AIH_SIPORT............ 0015
  +11B6  AIH_SISOCKTYPE........ 00000000
  +11BA  AIH_SIPROTOCOL........ 00000000
  +11BE  AIH_SIPORT............ 0000
        -- End of array   --
        AIH_DADDRS............
  +11C4  00000000  00000000  00000000  00000000  45937D41  00000002  00000000  00000000  00000000  43C3A04C  00000002  00000000

       ...  edited out part of output  ...
{{{{  ***** END OF OUTPUT ***** }}}}
```

SETPRINT command

Use the SETPRINT command to change the destination of subsequent IPCS command output. If the IPCSPRNT data set is allocated and being sent to a node, the output of future IPCS commands accumulates (but not displayed at the terminal) until you exit IPCS. When you exit IPCS, the IPCSPRNT data set is sent to the specified node.

Syntax

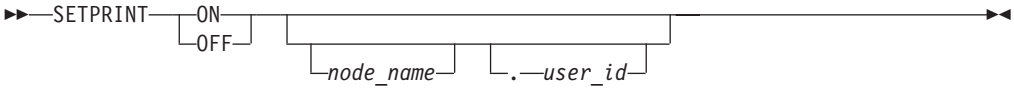

Parameters

ON Allocates the IPCSPRNT data set and issues the IPCS SETDEF PRINT command.

OFF
Frees the IPCSPRNT data set and issues the IPCS SETDEF NOPRINT command.

node_natme
The name of a TSO or VM system to which the output is sent.

user_id
The user ID on the TSO or VM system to which the output is sent.

Note: If *user_id* is specified, there must be a period but no space between *node_name* and *user_id*.

Sample output of the SETPRINT command

If the command completes successfully, there is no output for the SETPRINT command. The following examples are invalid invocations of the SETPRINT command.

Allocating IPCSPRNT when it is already allocated:

```
setprint on ralvms.testid
 IKJ56861I  FILE IPCSPRNT NOT UNALLOCATED, DATA SET IS OPEN
```

Freeing IPCSPRNT when it is already freed:

```
setprint off
 BLS21060I PRINT file not open
 IKJ56247I FILE IPCSPRNT NOT FREED, IS NOT ALLOCATED
```

SKMSG command

This topic describes the SKMSG command.

Use the SKMSG command to display the SKMSG fields.

Syntax

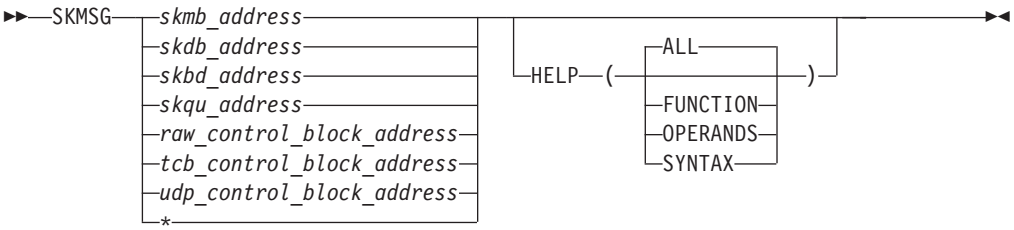

Parameters

***** To omit this positional parameter when using the HELP keyword.

skmb_address
The address of an SKMB control block or the symbol for the address.

skdb_address
The address of an SKDB control block or the symbol for the address.

skbd_address
The address of an SKBD control block or the symbol for the address.

skqu_address
The address of an SKQU control block or the symbol for the address.

raw_control_block_address
The address of a RAW control block or the symbol for the address.

tcb_control_block_address
The address of a TCB control block or the symbol for the address.

udp_control_block_address
The address of a UDP control block or the symbol for the address.

HELP
Display SKMSG usage and syntax information.

ALL
Displays help about the function, operands, and syntax information for the SKMSG command. ALL is the default.

FUNCTION
Display only function help information.

OPERANDS
Display only operands help information.

SYNTAX
Display only syntax help information.

Rule: If you specify multiple keywords from the set {ALL, FUNCTION, OPERANDS, SYNTAX}, all of the keywords that you specify are used.

Sample output of the SKMSG command

The following is a sample output of the SKMSG command:

```
SKMSG 15D4D5B8

SKDB at 15D4D5B8


Message 1
```

```
SKMB: 15D4D588
   +0000  id....... SKMB      next..... 00000000  prev..... 00000000
   +0008  tail..... 00000000  cont..... 15D55B08  flag..... 4000
   +0012  band..... 00        strx..... 16        hold..... 00000000
   +0018  datb..... 15D4D5B8  atch..... 00000000  rptr..... 15DE5A60
   +0024  wptr..... 15DE5AA8

SKDB: 15D4D5B8
   +0000  id....... SKDB      msgb..... 15D4D588  cver..... 00
   +0009  csrc..... 80        ctyp..... 40
   +000C  tokn..... 15E3F040  15E3F3E0  00001358  alet..... 00000000
   +001C  base..... 15DE5000  size..... 00001000  flag..... 00000000
   +0028  ref...... 00000005

Buffer: 15DE5000
   +0000  450005DC  24760000  4006DBF0  09433116  | ........ ..0.... |
   +0010  0943311A  0AB70866  481080F3  450C8271  | ...........3..b. |
   +0020  8010FFFE  DA3B0000  0101080A  3E456F23  | ..............?. |
   +0030  3E450C82  91A38897  D3C7E5D6  C297E8E3  | ...bjthpLGVOBpYT |
   +0040  D9F0E596  F2C989D9                      | R0Vo2IiR         |

SKMB: 15D55B08
   +0000  id....... SKMB      next..... 00000000  prev..... 00000000
   +0008  tail..... 00000000  cont..... 00000000  flag..... 0000
   +0012  band..... 00        strx..... 00        hold..... 00000000
   +0018  datb..... 15D55B38  atch..... 00000000  rptr..... 13ECA758
   +0024  wptr..... 13ECACB8

SKDB: 15D55B38
   +0000  id....... SKDB      msgb..... 15D55B08  cver..... 00
   +0009  csrc..... 40        ctyp..... 80
   +000C  tokn..... 15D41040  15D417F0  000003C7  alet..... 01FF0007
   +001C  base..... 13ECA000  size..... 00000CB8  flag..... 00800000
   +0028  ref...... 00000003
```

TCPHDR command

Use the TCPHDR command to display the TCP header fields.

Syntax

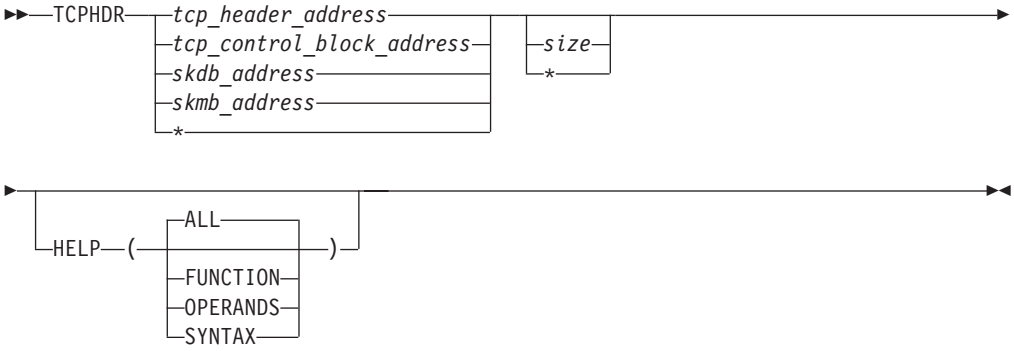

Parameters

***** To omit this positional parameter when using the HELP keyword.

tcp_header_address
The address of the TCP header or an IPCS symbol.

tcp_control_block_address
The address of a TCP/IP TCP control block or an IPCS symbol.

skdb_address
The address of an SKDB control block or an IPCS symbol.

skmb_address
The address of an SKMB control block or an IPCS symbol.

size
The amount of data to display. If the size is greater than the size of the header, the variable portion of the header (if it exists) is displayed. Must be one to three hexadecimal digits.

HELP
Display TCPHDR usage and syntax information.

ALL
Displays help about the function, operands, and syntax information for the TCPHDR command. ALL is the default.

FUNCTION
Display only function help information.

OPERANDS
Display only operands help information.

SYNTAX
Display only syntax help information.

Rule: If you specify multiple keywords from the set {ALL, FUNCTION, OPERANDS, SYNTAX}, all of the keywords that you specify are used.

Sample output of the TCPHDR command

The following is sample output of the TCPHDR command:

```
TCPHDR 7F522108

TCB at 7F522108

TCP Header at 7F5222D8

7F5222D8  04010402  7228DD16  7228DB82  50107FD8  | ..........b&."Q |
   +0010  00000000                                | ....            |

  Source Port         : 1025
  Destination Port    : 1026
  Sequence Number     : 1,915,280,662
  Ack Number          : 1,915,280,258
  Header Length       : 20
  Flags               : Ack
  Window Size         : 32728
  Checksum            : 0000
  Urgent Data Pointer : 0000
```

TOD command

Use the TOD command to format a hexadecimal time-of-day value into a readable date and time.

Syntax

Parameters

time_value
The time to be converted. *The time_value* can be specified as either 16 hexadecimal digits or as an address in a dump of an eight-byte STCK value. If less than 16 digits are specified, the value is padded on the right with zeros. If an address is specified, it must be followed by a period. If an address is less than eight hexadecimal digits, it is padded on the left with zeros.

time_zone
An offset for the time (the difference between local time and GMT). The *time_zone* can be specified either as a word or as a positive or negative decimal value. The recognized words are:

LOCAL
Time zone value of zero is used. This is the default.

GMT Greenwich Mean Time

EDT U.S. Eastern Daylight Time zone

EST U.S. Eastern Standard Time zone

CDT U.S. Central Daylight Time zone

CST U.S. Central Standard Time zone

MDT U.S. Mountain Daylight Time zone

MST U.S. Mountain Standard Time zone

PDT U.S. Pacific Daylight Time zone

PST U.S. Pacific Standard Time zone

Sample output of the TOD command

The following are sample outputs of the TOD command.

Sample output for STCK time-of-day with a time zone word:

```
Command ===> ip tod b214030791f3a92c,est

B2140307 91F3A92C : 1999/04/10 20:51:58.684986  TIMEZONE: 0000430E23400000
```

Sample output for an address in the dump where an STCK time-of-day value is located with a negative time zone offset:

```
Command ===> ip tod 11275d4.,-4

B24000E0 51900000 : 1999/05/16 05:36:37.632256  TIMEZONE: FFFFCA5B17000000
```

UDPHDR command

Use the UDPHDR command to display the UDP header fields.

Syntax

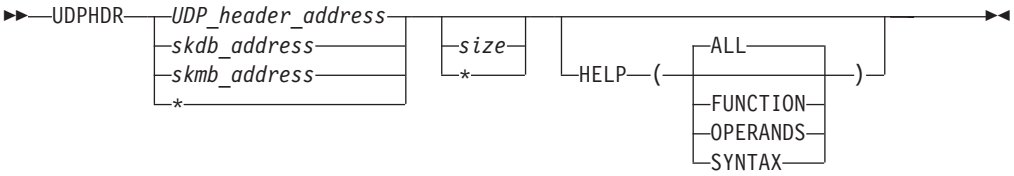

Parameters

***** To omit this positional parameter when using the HELP keyword.

UDP_header_address
The address of a UDP header or the symbol for the address.

Note: The UDP header has no version or identifier, so it is not possible to definitively recognize a UDP header given an address in storage. Therefore, this command formats the storage assuming it is a UDP header.

skdb_address
The address of an SKDB control block or the symbol for the address.

skmb_address
The address of an SKMB control block or the symbol for the address.

HELP
Display UDPHDR usage and syntax information.

ALL
Displays help about the function, operands, and syntax information for the UDPHDR command. ALL is the default.

FUNCTION
Display only function help information.

OPERANDS
Display only operands help information.

SYNTAX
Display only syntax help information.

Rule: If you specify multiple keywords from the set {ALL, FUNCTION, OPERANDS, SYNTAX}, all of the keywords that you specify are used.

Sample output of the UDPHDR command

The following is a sample output of the UDPHDR command:

```
UDPHDR 08D0A0D8

 UDP Header at 08D0A0EC

 08D0A0EC  040700A1  0033CD23                    | ...~....        |

   Source port         : 1031
   Destination port    : 161
   Datagram Length     : 51
   Checksum            : CD23
```

Installing TCP/IP IPCS subcommands by using the panel interface

To use the panel interface to the TCP/IP IPCS subcommands, you can either display the panels using an IPCS command or connect the TCP/IP ISPF panels to an existing ISPF panel. No additional installation steps are required to display the panels using an IPCS command. To connect the TCP/IP ISPF panels to an existing panel, find an existing panel where you wish to add TCP/IP as an option and modify the panel. Modify the panel by adding the TCP/IP option, which processes the following command:

```
PGM(BLSGSCMD) PARM(%EZBTCPEX) NEWAPPL(EZBD)
```

where `BLSGSCMD` is the IPCS command, `EZBTCPEX` is the TCP/IP REXX exec, and `EZBD` is the TCP/IP key list prefix.

You can also start the TCPIP panel interface by performing the following steps:

1. Log on to TSO.
2. Access IPCS to display the IPCS Primary Option menu. Figure 27 on page 310 shows an example of an IPCS Primary Option Menu.
3. Select option 2, "ANALYSIS - Analyze dump contents".
4. Select option 6, "COMPONENT - MVS component data".
5. Scroll down to "TCPIP TCP/IP Dump Analysis" and select it.

Entering TCP/IP IPCS subcommands

You can enter the TCP/IP IPCS subcommands as an IPCS command, either by using panels provided by TCP/IP or by using the IPCS batch facility.

You can invoke the TCP/IP IPCS panels in one of the following ways:

- Invoke the panel REXX exec as an IPCS subcommand. Follow the steps above for entering a TCP/IP IPCS subcommand using the IPCS Subcommand Entry panel and enter the command:

  ```
  EZBTCPEX
  ```

- Invoke the TCP/IP IPCS panels by selecting the option provided in the installation section above.

For either method, you should see the main menu for the TCP/IP IPCS commands shown in Figure 26 on page 310.

Select an option, and the panels prompt you for additional menu choices or input for the specific TCP/IP IPCS subcommand you select. After all input has been selected, the TCP/IP IPCS subcommand is invoked using the current default dump data set. If the dump data set is for Telnet, only the commands indicated "Available for Telnet" in the IPCS command list provide data. The commands not supported by Telnet return no data.

Restriction: For those TCP/IP IPCS subcommands that support storage addresses as input, a panel is displayed with a List field where you can enter multiple addresses for the subcommand. The List field only supports 31-bit storage addresses. 64-bit storage addresses are not supported. If you are interested in a specific 64-bit storage address related to a subcommand, invoke the subcommand without specifying any addresses in the List field. The resulting display includes the 64-bit storage addresses. You can then use the IPCS FIND command to locate the 64-bit storage addresses that you are interested in.

```
                          TCP/IP Analysis Menu
Command ===>
(C) Copyright IBM Corporation 1998,2001. All rights reserved.
Select one of the following.  Then press Enter.

    1.  General . . . - HEADER, MTABLE, STATE, TSDB, TSDX, TSEB
    2.  Protocol  . . - PROTOCOL, RAW, TCB, UDP
    3.  Configuration - CONFIG, CONNECTION, PROFILE, ROUTE
    4.  Resources   . - COUNTERS, LOCK, MAP, STORAGE, TIMER
    5.  Execution . . - DUAF, DUCB, TRACE
    6.  Interfaces  . - API, SOCKET, STREAM
    7.  Structures  . - HASH, TREE
    8.  Functions . . - FRCA, IPSEC, POLICY, TELNET, TTLS, VMCF,
                        XCF
    9.  Headers . . . - ICMPHDR, IPHDR, SKMSG, TCPHDR, UDPHDR
    10. Converters  . - ERRNO, SETPRINT, TOD
    11. Applications. - RESOLVER
```

Figure 26. Main menu for TCP/IP IPCS subcommands.

See z/OS Communications Server: New Function Summary for information regarding VTAM IPCS commands.

Steps for entering a TCP/IP IPCS subcommand

Enter a TCP/IP IPCS subcommand by using the IPCS Subcommand Entry panel.

Procedure

Follow these steps to enter a TCP/IP IPCS subcommand:

1. Log on to TSO.
2. Access IPCS to display the IPCS Primary Option Menu. Figure 27 shows an example of an IPCS Primary Option Menu.

```
------------------------  IPCS PRIMARY OPTION MENU  --------
OPTION   ===>
   0   DEFAULTS    - Specify default dump and options
   1   BROWSE      - Browse dump data set
   2   ANALYSIS    - Analyze dump contents
   3   UTILITY     - Perform utility functions
   4   COMMAND     - Enter IPCS subcommand or CLIST
   5   TCP/IP      - TCP/IP analysis commands
   6   NCP         - NCP analysis commands
   7   NMP         - NMP analysis commands
   8   INVENTORY   - Inventory of problem data
   9   SUBMIT      - Submit problem analysis job to batch
   T   TUTORIAL    - Learn how to use the IPCS dialog
   X   EXIT        - Terminate using log and list defaults

Enter END command to terminate IPCS dialog
```

Figure 27. IPCS primary option menu

3. Select option 4, COMMAND.
4. Type the TCP/IP IPCS subcommand. Figure 28 on page 311 shows the IPCS Subcommand Entry panel with a subcommand entered.

```
------------------------- IPCS Subcommand Entry --------------------------------
Enter a free-form IPCS subcommand or a CLIST or REXX exec invocation below:

===> tcpipcs help


----------------------- IPCS Subcommands and Abbreviations ---------------------
ADDDUMP            | DROPDUMP, DROP D  | LISTMAP, LMAP      | RUNCHAIN, RUN C
ANALYZE            | DROPMAP,  DROP M  | LISTSYM, LSYM      | SCAN
ARCHECK            | DROPSYM,  DROP S  | LISTUCB, LIST U    | SELECT
ASCBEXIT, ASCBX    | EQUATE,   EQU, EQ | LITERAL            | SETDEF,   SET D
ASMCHECK, ASMX     | FIND,     F       | LPAMAP             | STACK
CBFORMAT, CBF      | FINDMOD,  FMOD    | MERGE              | STATUS,   ST
CBSTAT             | FINDUCB,  FIND U  | NAME               | SUMMARY,  SUMM
CLOSE              | GTFTRACE, GTF     | NAMETOKN           | SYSTRACE
COPYDDIR           | INTEGER           | NOTE,    N         | TCBEXIT,  TCBX
COPYDUMP           | IPCS HELP, H      | OPEN               | WEBBEXIT, WEBBX
COPYTRC            | LIST,     L       | PROFILE, PROF      | WHERE,    W
CTRACE             | LISTDUMP, LDMP    | RENUM,   REN       |
```

Figure 28. IPCS subcommand entry panel with a TCP/IP IPCS subcommand entered

Step for using the batch option

Access IPCS commands by using the batch processing interface.

Procedure

Perform the following step:

Prepare the JCL data set. See z/OS MVS IPCS User's Guide and z/OS MVS IPCS Commands.
The following is a sample command (single command):

```
%TCPIPCS TELNET (* DETAIL)
```

Part 3. Diagnosing z/OS Communications Server components

© Copyright IBM Corp. 2000, 2015

Chapter 7. Diagnosing problems with the z/OS Load Balancing Advisor

The z/OS Load Balancing Advisor is a system that comprises outboard load balancers (LBs), an Advisor, and one or more Agents.

This topic discusses problem diagnosis of the Advisor and Agents and includes the following sections:

- "Diagnostic data"
- "Diagnosing Advisor and Agent problems" on page 316
- "Debug settings and corresponding syslogd priority levels" on page 320

Tip: For diagnosing problems with the load balancer, see the appropriate load balancer documentation.

Diagnostic data

You might need to collect multiple pieces of data to accurately diagnose problems. For example, the following might be useful:

- Console messages for Advisor and Agents
- Output from the MODIFY command for the Advisor and Agents
- syslogd log messages for Advisor and Agents (possibly including debug level trace)
- Advisor and Agent address space dumps and snap output
- Packet traces of Load Balancer data
- TCP/IP CTRACE of Agents and possibly the Advisor
- Netstat displays on TCP/IP stacks managed by Agents
- SNMP information

Guideline: syslogd does not have to be running to run the Advisor or Agents; however, syslogd is the only logging facility that either the Advisor or its Agents is capable of using. Useful diagnostic information might be lost if syslogd is not running before the Advisor or Agents are run.

The Advisor and Agent trigger address space dumps when certain unexpected error conditions are encountered. Both a CEEDUMP and address space snap output are produced and written to the data sets or files that are specified by the start procedure CEEDUMP and CEESNAP DD statements, respectively.

If the Advisor or Agent abnormally terminate (for example a 0Cx abend occurs), an unformatted SYSMDUMP is produced and written to the data set that is specified by the start procedure SYSMDUMP DD statement. If you override the Language Environment run-time option TERMTHDACT during the installation or start procedure, the SYSMDUMP might not be produced, or a CEEDUMP might be produced instead. Therefore, you should not override the TERMTHDACT run-time option. See z/OS Language Environment Programming Guide for more information about run-time options.

In other situations, the z/OS operator might need to dump the address space manually.

© Copyright IBM Corp. 2000, 2015

Packet trace data of Server/Application State Protocol (SASP) protocol messages that are sent between the Advisor and LBs might be needed. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for details about how to use the IP packet trace facility.

Restriction: When encrypting data, the packet trace data will be encrypted. Use MESSAGE level log messages instead.

The TCP/IP CTRACE trace of the Agents provides some information about data that has been collected from the TCP/IP stack for determining availability and desirability metrics. If the Agent is managing a CINET environment, a TCP/IP CTRACE might be needed in each TCP/IP stack. A TCP/IP CTRACE on the Advisor or Agent TCP/IP stack might also show data that is flowing between the Agents and Advisor. On the Agent TCP/IP stack, the SOCKET, INTERNET, and IOCTL CTRACE options are useful. On the Advisor TCP/IP stack, the INTERNET and SOCKET options are useful. See "Component trace" on page 45 for more information.

The following Netstat displays on stacks that are managed by Agents might be useful:

HOME
: Indicates which interfaces exist and which stack owns them

ALLCONN
: Indicates the listening TCP sockets and UDP end-points

SNMP information gives information similar to Netstat displays.

Diagnosing Advisor and Agent problems

This topic includes diagnostic information about Advisor and Agent problems.

Abends

Messages and error-related information are usually sent to the MVS system console when an abend on the Advisor or Agent occurs. Perform a dump of the error unless the symptoms already match a known problem.

Workload not distributed to a particular application

Use the following checklist to determine why workload is not being distributed to an application:

- Verify that the Advisor is running and that an Agent is running on the MVS system that contains the application. If they are not running, start the Advisor or Agent.
- If you are using sysplex subplexing, verify that the Advisor and Agents are in the same subplex. If there are multiple TCP/IP stacks in a subplex, ensure the IP addresses used by the Advisor and Agents are DVIPAs defined within a VIPARANGE statement on each of the stacks in the subplex. Review the syslog for the Advisor and Agents for messages indicating what subplex had been used. Each subplex must have an active Advisor associated with it, and each subplex in a z/OS system must have an Agent associated with it.
- Issue display commands on the Advisor to determine whether any LBs have registered the application. Verify that the LB is connected to the Advisor. If you are using sysplex subplexing, ensure that the load balancers have connectivity to the Advisor's subplex.

- Verify that the Advisor's `lb_id_list` statement includes the IP address of the LB in question if not using AT-TLS.
- Verify that the IP address and protocol of the member on the LB match the IP address and protocol of the application. If the IP addresses or protocols do not match, correct the definition at the LB.
- Verify that the Advisor's `agent_id_list` statement contains the IP address and port that the Agent is bound to on the system where the application exists, if not using AT-TLS. If it does not match and you are not using AT-TLS, correct the `agent_id_list` statement on the Advisor or the `advisor_id` statement on the Agent.
- Verify that network connectivity exists between the Advisor and the Agent in question. Unexpected loss of network connectivity between the two should result in an immediate action console message and related messages in the Agent and Advisor log. Issue NETSTAT CONN or NETSTAT ALLCONN commands on the Advisor system to see which Agents have connections to the Advisor, and by omission, which do not. Issue the MODIFY DISPLAY command on the Agents in question to verify that the connection to the Advisor is still active. The DISC flag is shown on the MODIFY DISPLAY command when the Agent is not connected to an Advisor. Correct the underlying network connectivity problem. For more information, see Chapter 4, "Diagnosing network connectivity problems," on page 25.
- If using AT-TLS with SERVAUTH access control checks to validate connections between the Advisor, Agents, and external load balancers, see Chapter 28, "Diagnosing Application Transparent Transport Layer Security (AT-TLS)," on page 707. In addition, ensure that the SERVAUTH class is active. Ensure that the EZB.LBA.LBACCESS.sysname.tcpsysplexgroupname resource profile is defined and that the user ID associated with the external load balancer has READ access to it. Ensure that the EZB.LBA.AGENTACCESS.sysname.tcpsysplexgroupname resource profile is defined and that the Agents have READ access to it. On the system console where the Advisor is running, look for message EZD1280I which indicates that a connection attempt using AT-TLS failed. This message has specific reason codes which indicate the reason for the failure.
- Issue display commands on the Advisor and Agent in question to verify that the application is available and enabled (not quiesced). Start the application or enable the application using the Agent MODIFY ENABLE command.
- Check the log file for ERROR or WARNING messages and take the appropriate corrective action. If ERROR and WARNING level log messages are not enabled, enable them and recheck the log file later.
- Verify that the LB has connectivity to the IP address of the member in question.

Workload not distributed as expected

Use the following checklist to diagnose workload distribution problems:

- Verify that the Advisor's update interval value is not inordinately large. The Advisor must wait at least two update intervals before beginning to receive enough data to properly calculate weights when an application becomes available or when an Agent is started.

  Allow at least three update intervals to expire after an application is started before re-examining the distribution of workload. If workload is occasionally being sent to overloaded applications or systems, adjust the update_interval downward so workload distribution can react more quickly to the pace of new workload requests.
- Periodically issue display commands on the Advisor to check the weights of members within the group in question. Determine whether the weights are

consistent with the expected behavior. If the weights are not consistent with expected behavior, see z/OS Communications Server: IP System Administrator's Commands for more information about how to analyze the member weights; if all releases in the sysplex are not V1R9 or above, note the restrictions and limitations described in this section. If these are consistent with the expected behavior, investigate the problem at the LB. For more information, see z/OS Communications Server: New Function Summary.

- Verify that the Advisor's `agent_id_list` value contains the IP addresses and ports that each Agent is bound to on the MVS systems where the application exists, if not using AT-TLS. If it does not match and you are not using AT-TLS, correct the `agent_id_list` statement on the Advisor or the `advisor_id` statement on the Agent.
- If using AT-TLS with SERVAUTH access control checks to validate connections between the Advisor, Agents, and external load balancers, see Chapter 28, “Diagnosing Application Transparent Transport Layer Security (AT-TLS),” on page 707. In addition, ensure that the SERVAUTH class is active. Ensure that the EZB.LBA.LBACCESS.sysname.tcpsysplexgroupname resource profile is defined and that the user ID associated with the external load balancer has READ access to it. Ensure that the EZB.LBA.AGENTACCESS.sysname.tcpsysplexgroupname resource profile is defined and that the Agents have READ access to it. On the system console where the Advisor is running, look for message EZD1280I which indicates that a connection attempt using AT-TLS failed. This message has specific reason codes which indicate the reason for the failure.
- Issue display commands at the Advisor to make sure that members of the group in question are not unexpectedly quiesced or unexpectedly unavailable (AVAIL status is NO).

  If AVAIL is NO because sysplex problem detection and recovery issued message EZD1973E then the TCP/IP stack must either be restarted or at least 10 minutes must pass from the last occurrence of the problem (for example, an abend). See the Sysplex problem detection and recovery section in z/OS Communications Server: IP Configuration Guide for more information.
- Issue display commands at the Advisor for all system-level members in the sysplex to verify that the MVS systems have the expected residual capacity.
- Check the log file for ERROR or WARNING messages and take the appropriate corrective action. If ERROR and WARNING level log messages are not enabled, enable them and recheck the log file later.

Advisor or Agent appears to be hung

Verify that the Agent or Advisor is actually hung by issuing a MODIFY *procname*,DISPLAY,DEBUG command. If no response is received, then attempts to stop (not cancel) the application. If the application does not terminate, the application is hung. If the hang occurred while DEBUG-level Advisor or Agent trace was in effect, collect the following problem documentation and call IBM Service.

- Take an SVC dump of the Agent or Advisor address space (depending on which application is hung) and of the OMVS address space including its data spaces.
- Capture the MVS console messages.
- Capture the application (Agent or Advisor) log messages written to syslogd.

If DEBUG-level trace was not in effect at the time, turn on DEBUG-level Advisor or Agent trace, reproduce the problem, collect the problem documentation previously mentioned, and call IBM Service.

Group names in displays are indecipherable

When LBs define group names, the names are coded in UTF-8 format. This character set is a superset of the EBCDIC character set; therefore, not all characters are translatable to EBCDIC. Rename the group names in the LBs to use characters limited to the ASCII character set.

Load balancer connection terminates unexpectedly

Check the following:

- Verify the load balancer administrator has not shut down the load balancer.
- Verify that TCP/IP connectivity still exists between the load balancer and the Advisor (for example, from the Advisor host, ping the address of the load balancer).
- Check the Advisor's log file for ERROR or WARNING messages and take the appropriate corrective action. If you see an ERROR message indicating a send() operation failed with "errno = EDC8102I Operation would block" you might have too many groups or members registering from the load balancer. Increase the TCPSENDBFRSIZE parameter of the TCPCONFIG PROFILE.TCPIP statement, or register fewer groups and members from the load balancer, and then try the operation again. If ERROR and WARNING level log messages are not enabled, enable them, repeat the operation, and recheck the log file again.

  Tip: Keep in mind that the Advisor has an internal maximum message size of 128K bytes. If this limit is exceeded, the connection is closed and an error message is logged stating that the message is too large and was not received.
- Check the load balancer for errors.

Agent-Advisor connection terminates unexpectedly

Check the following:

- Verify that the Agent's MVS operator has not shut down the Agent.
- Verify that TCP/IP connectivity still exists between the Agent and the Advisor.
- Check the Advisor's log file for ERROR or WARNING messages and take the appropriate corrective action. If you see an ERROR message indicating a send() operation failed with `errno = EDC8102I Operation would block`, you might have too many groups or members registered that belong to the same Agent. Increase the TCPSENDBFRSIZE parameter of the TCPCONFIG PROFILE.TCPIP statement, or register fewer groups and members belonging to the Agent. Then try the operation again.

  Tip: Keep in mind that the Advisor and Agent have an internal maximum message size of 128KB. If this limit is exceeded, the connection is closed and an error message is logged, which states that the message is too big and was not received.

  If ERROR and WARNING level log messages are not enabled, enable them, repeat the operation, and recheck the log file again.
- Check the Agent's log for errors.
- Connectivity can be dropped between the Advisor and Agents if processing slows down too much (due to lots of registered members and/or high debug levels) and the update_interval is configured too low. Using AT-TLS could increase the possibility of this happening. Try increasing the update_interval.

Automatic restart manager (ARM) registration failure

Failure of the Advisor or Agent to properly register with ARM is indicated by a warning-level message written to the log file. This log message is a result of the IXCARM call failing with the return code and reason codes indicated in the log message. See z/OS MVS Programming: Sysplex Services Reference for information about interpreting the IXCARM return code and reason code.

One of the common causes of failure is the lack of a security profile. See the EZARACF sample for instructions on how to add an ARM security profile for the application.

When you are using sysplex subplexes, do the following:

- Define an ARM policy with the TARGET_SYSTEM keyword to indicate which systems the element can be restarted on to ensure that the application is restarted only on a system that is in the same subplex.
- Restart the Advisor and Agent on a VTAM system that has been started with an XCFGRPID start option that corresponds with the vv portion of the *sysplex_group_name* in the Advisor and Agent configuration files, and has an available TCP/IP stack with a GLOBALCONFIG XCFGRPID parameter that corresponds with the tt portion of the *sysplex_group_name* in the Advisor and Agent configuration files.
- If there are multiple TCP/IP stacks in a subplex, ensure the IP addresses used by the Advisor and Agents are DVIPAs defined within a VIPARANGE statement on each of the stacks in the subplex.

Debug settings and corresponding syslogd priority levels

Table 16 summarizes the available debug levels and their associated syslogd priority levels.

Table 16. Available debug levels and associated syslogd priority levels

Logging category/Level	Description
None — 0	No messages of any kind are sent to the logging file after initialization is complete.
ERROR — 1	Error messages indicate something that requires attention. Messages at this level could be fatal (terminating) or could indicate that an important part of the workload advising system is not working properly. This information is logged at the syslogd ERROR priority level.
WARNING — 2	Warning messages indicate that an error has occurred, but it is not severe enough to warrant an ERROR. Corrective action might be necessary because the Advisor or Agent might not be behaving as intended. This information is logged at the syslogd WARNING priority level.

Table 16. Available debug levels and associated syslogd priority levels (continued)

Logging category/Level	Description
EVENT — 4	Event messages are logged for things that happen periodically, like operator commands, UNIX signals, timer pops, receipt of a network message, and so on. This information is logged at the syslogd NOTICE priority level.
INFO — 8	Informational messages are sent to the logging file. These messages do not require corrective action. This information is logged at the syslogd INFO priority level.
MESSAGE — 16	Message messages concern the detailed contents of message packets that are sent between the Advisor and LB, or between the Advisor and Agent. These can be used to assist debugging Advisor/LB and Advisor/Agent communications. This information is logged at the syslogd DEBUG priority level. This level is intended for IBM service use only.
COLLECTION — 32	Collection messages concern the details of collecting and manipulating the data that forms the basis of weight calculations. This information is logged at the syslogd DEBUG priority level. **Restriction:** COLLECTION is only used by the Agent. This level is intended for IBM service use only.
DEBUG — 64	Debug messages are intended for Development or Service and give detail that customers would not normally want. The intention of this level of message is to provide information that is useful in debugging code, logic, or timing errors. This information is logged at the syslogd DEBUG priority level. This level is intended for IBM service use only.
TRACE — 128	Trace messages are intended for Development or Service to track code processing (footprints). This information is logged at the syslogd DEBUG priority level. This level is intended for IBM service use only.

Chapter 8. Diagnosing problems with the automated domain name registration application (ADNR)

The automated domain name registration (ADNR) application is a function that dynamically updates name servers with information about sysplex resources in near real time. The DNS names managed by ADNR can be names that represent all instances of an application within the sysplex, names that represent a specific instance of an application within the sysplex, names that represent the entire sysplex, or names that represent individual systems within the sysplex. ADNR communicates with the z/OS Load Balancing Advisor, (specifically the Advisor application), which architecturally is a Global Workload Manager (GWM) according to the Server/Application State Protocol (SASP). The Advisor application from the z/OS Load Balancing Advisor is the only GWM with which ADNR is designed to interact. All references to a GWM in this topic see the Advisor application of the z/OS Load Balancing Advisor.

Diagnostic data

You might need to collect multiple pieces of data to accurately diagnose problems, such as the following:

- Console messages for ADNR
- syslogd log messages for ADNR (possibly including DEBUG - 64 level trace)
- Name server log data for the name servers that are managed by ADNR
- ADNR address space dump and snap output
- SYSTCPIP CTRACE for the TCP/IP stack where ADNR and the GWM are running
- Packet traces of GWM data
- Netstat displays for the connection between ADNR and the GWM
- A listing of the zone data from the managed name server or name servers
- z/OS Load Balancing Advisor log data and displays. See Chapter 7, “Diagnosing problems with the z/OS Load Balancing Advisor,” on page 315 for information about the Load Balancing Advisor.

Note: syslogd is the only logging facility that ADNR uses. Useful diagnostic information might be lost if syslogd is not running before you run ADNR.

ADNR triggers address space dumps when certain unexpected error conditions are encountered while communicating with a GWM. Just after connecting to the GWM, ADNR enters a negotiation phase. During negotiation, a series of architected SASP requests are sent to the GWM; for each request, an architected SASP reply is received from the GWM. If the negotiation does not successfully complete, ADNR closes the connection, increases the logging level, establishes a new connection to the GWM, and retries the negotiation.

If the negotiation fails a second time, ADNR dumps its address space; the dump title header is ADNR Dump - Neg Failed and the logging level is restored to its original configured value and the connection is closed. Retries continue at 1 minute intervals using the configured logging level.

© Copyright IBM Corp. 2000, 2015

After completing the negotiation phase, the GWM might send an unsolicited SendWeights message to ADNR; the message contains information about the changed state of resources that the ADNR application registered to the GWM during negotiation. If the SendWeights message contains an architectural error, ADNR closes the connection, increases the logging level, establishes a new connection to the GWM, and completes negotiation with the GWM. If the next SendWeights message received from the GWM contains an error, ADNR dumps its address space; the dump title header is as follows:

```
ADNR Dump - Rcv Failed
```

The logging level is restored to its original configured value and the connection is closed. Retries continue at 1 minute intervals using the configured logging level.

These types of errors generally occur because incompatible levels of maintenance are applied to the GWM and ADNR. After being started, ADNR dumps its address space only once when these types of errors are detected. For further diagnosis, collect the ADNR log and address space dump. Review the log to determine the type of error that occurred. Review the PTF requirements of any recently installed PTFs. If you cannot correct the problem with additional maintenance or a configuration change, then contact IBM Service.

For other types of errors, both a CEEDUMP and address space snap output might be produced and written to the data sets or files that are specified by the start procedure CEEDUMP and CEESNAP DD statements, respectively.

If ADNR abnormally terminates (for example an 0Cx abend occurs), then an unformatted SYSMDUMP is written to the data set that is specified by the start procedure SYSMDUMP DD statement. If you override the Language Environment run-time option TERMTHDACT during the installation or start procedure, then the SYSMDUMP may not be produced, or a CEEDUMP may be produced instead. Therefore, you should not override the TERMTHDACT run-time option. See z/OS Language Environment Programming Guide for more information about Language Environment runtime options.

In other situations, the z/OS operator might need to dump the ADNR address space manually. See z/OS MVS Diagnosis: Tools and Service Aids for more information about obtaining a dump.

SYSTCPDA CTRACE (packet trace) data of the SASP protocol messages sent between ADNR and GWM may be needed. See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45 for details about how to use the IP packet trace facility.

Restriction: When encrypting data, the packet trace data will be encrypted. Use MESSAGE level log messages instead.

A SYSTCPIP component trace on the TCP/IP stacks used by ADNR and its associated GWM shows data that is flowing between them. Start the trace by specifying OPTIONS=(PFS,TCP,UDP,INTERNET),JOBNAME=(*server*) on both stacks, where the *server* value is the ADNR or GWM address space (or both names separated by a comma if they are using the same stack). See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45 for more information.

You can dump the contents of the DNS zones managed by ADNR by issuing the domain information grouper (**dig**) command from z/OS UNIX:

```
dig @server zone axfr -p port -k key_file 2>zone_xfer.err 1>zone_xfer.out
```

where

server
The server IP address or host name which contains the zone managed by ADNR

zone
The zone being managed by ADNR whose contents are being dumped

port
Optionally specify the port number which the DNS server is listening on for queries

key_file
Optionally specify the key file that is used to sign transactions for this zone

zone_xfer
Redirect the stdout and stderr file streams of the command to two distinct z/OS UNIX files.

See z/OS Communications Server: IP System Administrator's Commands for more information about the **dig** command.

The following Netstat command displays stacks that have affinity with ADNR:

```
ALLConn/-a, COnn/-c
```

This command is used to determine whether there is an active connection between ADNR and the GWM. ALLConn/-a displays information for all TCP connections and UDP sockets, including some recently closed ones. COnn/-c displays information about each active TCP connection and UDP socket. See z/OS Communications Server: IP System Administrator's Commands, Monitoring the TCP/IP network, Netstat section for information about using Netstat commands.

Diagnosing ADNR problems

This topic includes diagnostic information about ADNR problems.

Abends

Messages and error-related information are usually sent to the MVS system console when an abend on ADNR occurs. Perform a dump of the error unless the symptoms already match a known problem.

ADNR fails to initialize

Problems with the configuration file are the most common cause for ADNR failure during initialization. This class of problems is identified by a console message. However, failure to give the ADNR load module proper APF authorization will not result in an ADNR termination message on the console or in the syslog. In this particular case, the failure message is sent to the SYSOUT data set. If the sample ADNR started procedure is used, then this output appears in the ADNR job log.

ADNR not communicating with the Global Workload Manager

ADNR communicates with only one GWM. Use the following information to diagnose why ADNR fails to communicate with the GWM.

Restriction: ADNR supports only the z/OS Load Balancing Advisor application as the GWM.

- Verify that ADNR is running. If it is not running, then start ADNR.
- Verify that the GWM is running. If it is not running, then start the GWM. See Chapter 7, “Diagnosing problems with the z/OS Load Balancing Advisor,” on page 315 for more on z/OS Load Balancing Advisor problems.
- Verify network connectivity exists between the GWM and ADNR.
  - See Chapter 4, “Diagnosing network connectivity problems,” on page 25 for guidance on the diagnosis process for network connectivity problems.
  - Issue display commands on the GWM (by using the MODIFY command) to determine whether ADNR is connected to the GWM and has registered its group and member data. For more information about the MODIFY command see z/OS Communications Server: IP System Administrator's Commands. If not using AT-TLS, verify that the z/OS Load Balancing Advisor's `lb_id_list` statement includes the IP address on the `host_connection_addr` parameter in the gwm statement in the ADNR configuration file. The eventual action message, EZD1272E will persist on the console if communication with a GWM does not exist.
  - If using AT-TLS with SERVAUTH access control checks to validate connections between the Advisor and ADNR, see Chapter 28, “Diagnosing Application Transparent Transport Layer Security (AT-TLS),” on page 707. In addition, ensure that the SERVAUTH class is active. Ensure that the EZB.LBA.LBACCESS.sysname.tcpsysplexgroupname resource profile is defined and that the load balancer and ADNR have READ access to it. On the system console where the Advisor is running, look for message EZD1280I which indicates that a connection attempt using AT-TLS failed. This message has specific reason codes which indicate the reason for the failure.
  - If you are using sysplex subplexing, ensure that the ADNR application does the following:
    - Has connectivity to the Advisor's subplex.
    - Is using one of the TCP/IP stacks in this subplex; when in a common INET (CINET) environment with multiple TCP/IP stacks on one MVS system, either by establishing affinity to one of the stacks in the subplex or by binding to a VIPA that is only defined on stacks that are in that subplex. See the adnrproc.sample for an example of the JCL to establish affinity.
  - Verify that network connectivity exists between ADNR and the GWM in question. Issue Netstat COnn/-c or Netstat ALLConn/-a commands on the ADNR system to see whether a connection exists between ADNR and the GWM. Correct the underlying network connectivity problem. For guidance on using Netstat commands see z/OS Communications Server: IP System Administrator's Commands.
  - Issue a display command on ADNR to determine whether there are indications that any groups and members are known to exist within the sysplex. For more information about the MODIFY command see z/OS Communications Server: IP System Administrator's Commands.
- Check the syslogd output file where ADNR writes its log data for ERROR or WARNING messages and take the appropriate corrective action. If ERROR and WARNING level log message are not enabled, then enable them and recheck the log file later. The current ADNR debug level can be displayed by issuing the MODIFY procname,DISPLAY,DEBUG command at the MVS console. The debug level can be dynamically changed by issuing the MODIFY *procname*,DEBUG,LEVEL=*n* command at the MVS console. The *procname* is the JCL procedure name for ADNR and *n* is the new debug level. For more information about the MODIFY command, see z/OS Communications Server: IP System Administrator's Commands.

Automatic restart manager (ARM) registration failure

Failure of ADNR to properly register with ARM is indicated by a warning-level message written to the log file. This log message is a result of the IXCARM call failing with the return code and reason codes indicated in the log message. See z/OS MVS Programming: Sysplex Services Reference for information about interpreting the IXCARM return code and reason code.

One of the common causes of failure is the lack of a security profile. See the EZARACF sample for instructions on how to add an ARM security profile for the application. Ensure that each instance of ADNR is configured to use an ARM element name that is unique within the sysplex. Use the *arm_element_suffix* configuration statement to specify a unique suffix for the element name.

ADNR not updating zones in a DNS server

Use the following information to determine why changes to host names are not being updated in the DNS server zone being managed by ADNR:

- Verify that ADNR is running. If it is not running then start ADNR.
- If message EZD1278E or EZD1257I has been issued, see "Diagnosing unresponsive zones" on page 328.
- Issue display commands with the ADNR MODIFY command to determine whether it is connected to the GWM and has registered its group and member data. If not using AT-TLS, verify that the z/OS Load Balancing Advisor's lb_id_list value includes the IP address of ADNR specified with the host_connection_addr keyword of the gwm statement.

  If using AT-TLS with SERVAUTH access control checks to validate connections between the Advisor and ADNR, see Chapter 28, "Diagnosing Application Transparent Transport Layer Security (AT-TLS)," on page 707. In addition, ensure that the SERVAUTH class is active. Ensure that the EZB.LBA.LBACCESS.sysname.tcpsysplexgroupname resource profile is defined and that the load balancer and ADNR have READ access to it. On the system console where the Advisor is running, look for message EZD1280I which indicates that a connection attempt using AT-TLS failed. This message has specific reason codes which indicate the reason for the failure.
- Check the log file for any ADNR ERROR or WARNING messages and take the appropriate corrective action. If ERROR and WARNING level log messages are not enabled, then enable them and recheck the log file later. For more information about the ADNR display commands see z/OS Communications Server: IP System Administrator's Commands.
- Ensure that the ADNR configuration is not changed when ADNR is not active. Removing a dns statement or zone parameter while ADNR is not active causes ADNR to lose control of the information in that name server's zones. The information in this case is considered to be orphaned.

DNS name servers managed by ADNR contain incorrect or outdated data

Use the following information to determine why zones being managed by ADNR contain incorrect or outdated information:

- Verify that ADNR is communicating with the GWM and its managed name servers.
- Verify that ADNR is able to communicate with the managed DNS name server's zones.

- Ensure that the ADNR configuration file is not changed when ADNR is not active. Removing a dns statement or zone parameter while ADNR is not active causes ADNR to lose control of the information in that name server's zones. The information in this case is considered orphaned and goes stale. These types of configuration file changes should be made while ADNR is active and applied by using the MODIFY *procname*,REFRESH command to avoid orphaned data in the name server; ADNR deletes the information in the name server that is associated with the removed dns statement or zone parameter.
- Verify that the zones in the name servers managed by ADNR have not been updated by any entity other than ADNR. This includes manual updates to the zone data files, updates from DHCP servers, or other nsupdate clients. Failure to abide by this restriction can result in lost DNS records and ADNR zone update failures.
- Verify that the update interval of the GWM is not longer than you expect. The GWM update interval dictates how frequently ADNR receives data from the GWM and consequently, how frequently the managed name servers are updated with that data. Lower the update interval on the GWM if you need the managed name servers to have data that more closely follows the actual availability status of the sysplex resources they represent. ADNR waits a certain time after ADNR initialization and after a dynamic update to ensure that all of the sysplex data has been reported before attempting to update its managed name servers. When the GWM is the z/OS Load Balancing Advisor (Advisor) application, the Advisor's update_interval statement determines the time that ADNR waits; specifically two times the update_interval received from the GWM. If the ttl keyword under a zone parameter of the ADNR DNS statement is defaulted to use the value from the GWM's update_interval statement, then that value is used as the time-to-live value for the DNS resource records for that zone
- If the ttl value for a zone is defaulted as described in the previous bullet, ensure the resource records in the name server for that zone reflect this value.

Diagnosing unresponsive zones

Messages EZD1278E and EZD1257I indicate when a zone is unresponsive and identify an unresponsive zone. An unresponsive zone does not accept updates or queries for information from ADNR. Unresponsive zones cause other symptoms, such as zones in a name server not being updated at all, or failure of the zone to contain up-to-date information regarding the status of resources in the sysplex. Use the following information to determine why a zone is not responsive.

- Issue display commands through the ADNR MODIFY command to determine whether the name server is responding. A name server managed by ADNR can be comprised of one or more zones. A MODIFY *procname*,DISPLAY,DNS,DETAIL command shows a count of the number of zones defined under a dns statement and a count of the number of zones under that dns statement that are active. When all of the zones managed by a DNS server controlled by ADNR are not responding then the DNS server is considered dead. ADNR makes periodic probes to determine whether the zones for a dead server respond positively.
- Verify that the DNS server being used supports RFC 2136, *Dynamic Updates in the Domain Name System (DNS UPDATE)*). Review your DNS server's documentation.
- Verify the DNS name server is running and responsive by issuing the **dig** or **nsupdate** command for the zone. If it is not running then start the DNS server. See z/OS Communications Server: IP System Administrator's Commands, Querying and administrating a Domain Name System (DNS), for guidance on using the **nslookup**, **dig** and **nsupdate** commnands.

- Verify that network connectivity exists for the DNS server as it must be listening on the IP address and port number specified by the server parameter value (IP address) of the `dns` statement in the ADNR configuration. Unexpected loss of network connectivity for the DNS server will result in a console message and related messages in the DNS log. Correct the underlying network connectivity problem. For guidance on using Netstat commands see z/OS Communications Server: IP System Administrator's Commands.
- Review your firewall's log files to verify a firewall is not blocking communications between the system where ADNR resides and the name server on the port where the name server is listening for queries.
- Verify the name server being used is listening at the IP address and port that is specified by the dns_id parameter of the dns statement in the ADNR configuration file. For DNS BIND9, the IP addresses and ports the DNS server will listen on may be specified by the listen-on and listen-on-v6 DNS option statements.
- Verify that the name server IP address, optional port, zone domain suffix names, and optional Transaction Signature (TSIG) keys are correctly specified in the ADNR configuration file.
- Verify that the DNS name server specified in the ADNR dns statement actually manages the zone specified by the domain_suffix parameter of the ADNR dns statement and is the authoritative, primary master name server for the zone. For DNS BIND9 on the name server's zone configuration statement, the type master option is used to specify that the server is an authoritative master. See BIND 9-based domain name system (DNS), z/OS Communications Server: IP Configuration Reference, configuration file statements, for guidance on coding the name server's configuration. The name server managing this zone must be configured for the specified zone before ADNR can add DNS records to it. ADNR cannot dynamically create a zone in a name server. It can only add records to a zone that already exists.
- Verify that ADNR has the authority to manage the DNS resource records contained in this zone including the authority to request and receive zone transfers and perform dynamic updates. See Automated Domain Name Registration in z/OS Communications Server: IP Configuration Guide, for guidance on authorizing ADNR.
- Verify that the transaction security (TSIG) keys represented in the update and transfer keys (if specified in the ADNR configuration file) match those specified in the DNS name server for the zones ADNR is managing.
- Verify that the name server is configured with the same key names that ADNR is configured to use for the zone. Even if the name server configuration does not require a key to update or transfer an ADNR managed zone, the keys must at least be configured to the name server if ADNR is configured to use a key for that zone. If your security policies do not require you to use an update or a transfer key, they should be removed from the ADNR configuration, otherwise, the keys should be configured to the name server and used to restrict which entities are allowed to update the zone and request zone transfers.
- Verify that the name server's working directory did not run out of disk space. ADNR makes dynamic updates to name severs. Many name server implementations require that dynamic updates be written to disk. If a name server is unable to do this, the dynamic updates from ADNR will fail causing the zone to go unresponsive. In this case, the zone emerges from the unresponsive state spontaneously, but again returns to the unresponsive state. This cycle will repeat until the storage problem on the name server host is corrected

- Verify that the zone specified on the zone_label keyword of the ADNR dns configuration statement is not a DNSSEC signed zone. ADNR does not support the use of zones signed by DNSSEC.
- Verify that OMVS has not run out of file descriptors. See the DISPLAY OMVS command in z/OS MVS System Commands for information about how to make this determination.

ADNR appears to be hung

Verify that ADNR is actually hung by first issuing a MODIFY procname,DISPLAY,GWM command. If no response is received then attempt to stop (not cancel) the application. See z/OS MVS System Commands, STOP command subsection for more information about stopping an address space. If the application does not terminate, then the application is hung. If the hang occurred while the debug-level ADNR trace was in effect, then collect the following problem documentation and call IBM Service:

- Take an SVC dump of the ADNR address space
- Take an SVC dump of TCP/IP address space
- Take an SVC dump of the OMVS address space including its data spaces
- Capture the MVS console messages
- Capture the ADNR log messages written to syslogd

If the ADNR debug-level trace was not in effect at the time, then turn on the debug-level ADNR trace, reproduce the problem, collect the problem documentation, and call IBM Service.

ADNR connection to the GWM terminates unexpectedly

Determine whether the GWN is available, that TCP/IP connectivity still exists between ADNR and the GWM, and check for error and warning messages.

Procedure

Check the following:

- Verify that the load balancing administrator has not shut down the GWM advising ADNR.
- Verify that TCP/IP connectivity still exists between ADNR and the GWM (for example, from the ADNR host, ping the address of the GWM). See Monitoring the TCP/IP network, Ping subsection, in z/OS Communications Server: IP System Administrator's Commands, for further information about ping.
- Check ADNR's log file for ERROR or WARNING messages and take the appropriate corrective action. If ERROR and WARNING debug level log messages are not enabled, then enable them, repeat the operation, and recheck to log file again. See Modify command -- Automated Domain Name Registration in z/OS Communications Server: IP System Administrator's Commands, for further information about enabling ADNR's debug levels.
- Check the GWM for errors.

Debug settings

The value specified by the ADNR `debug_level` configuration option determines the ADNR logging levels. See z/OS Communications Server: IP Configuration Reference, Automated Domain Name Registration, Automated Domain Name Registration configuration file section for more on ADNR logging levels. The

values may be added together to trace multiple logging categories. See z/OS Communications Server: IP System Administrator's Commands, Operator commands and system administration, Modify command, Modify command -- Automated Domain Name Registration subsection for further information about displaying and changing ADNR's debug levels.

- When a problem occurs communicating with a DNS zone, then specifying a debug level of COLLECTION -32 causes ADNR to log the Nsupdate and Dig commands, and responses against the DNS zone. This data is the exact Nsupdate and Dig client commands and any associated responses
- When ADNR is not able to communicate with a GWM, then specifying a debug level of MESSAGE record (16) causes ADNR to log the SASP flows.

Chapter 9. Diagnosing IKE daemon problems

This topic describes how to diagnose IKE daemon problems, and contains the following subtopics:

- "Overview of diagnosing IKE daemon problems"
- "Diagnosing IKE daemon problems" on page 334
- "IKE daemon debug information" on page 350
- "TCP/IP services component trace for the IKE daemon" on page 351
- "Steps for enabling the CTRACE at IKE daemon startup" on page 354

Overview of diagnosing IKE daemon problems

This section provides overview information about the z/OS Internet Key Exchange (IKE) daemon and its functions.

The IKE daemon manages dynamic IPSec tunnels. The IKE daemon is not involved in the filtering, encapsulation, or decapsulation of packets. The IKE daemon is not required for the configuration or use of IP filters.

The critical elements of IP security are security associations (SAs); specifically the information that they provide about the partners of a secure communications channel, and the cryptographic algorithms and keys to be used. The z/OS IKE daemon supports two versions of the Internet Key Exchange: IKE version 1 (IKEv1) and IKE version 2 (IKEv2). The Internet Security Association Key Management protocol (ISAKMP) provides a framework for exchanging messages to automate the negotiation of security associations. The IKEv1 protocol is a hybrid protocol that conforms to the ISAKMP framework and implements a subset of the Oakley and SKEME protocols to negotiate SAs and provide authenticated keying material for SAs in a protected manner. The IKEv2 protocol is very similar to the IKEv1 protocol, in that it also negotiates SAs and provides authenticated keying material for SAs in a protected manner

The z/OS IKE daemon implements the IKE protocol to dynamically establish SAs with peer daemons that also support these protocols. In the sections that follow, a peer daemon might be referred to as an ISAKMP server or ISAKMP peer. Also, the z/OS IKE daemon might be referred to as the IKE daemon or IKED.

The IKE daemon establishes SAs within the guidelines of internet protocol security (IP security) policy. IP security policies are defined in one or more local files that are read by the Policy Agent. The IKE daemon obtains IP security policies from the Policy Agent using the Policy API (PAPI). See z/OS Communications Server: IP Configuration Guide for more information about configuring and starting Policy Agent, as well as defining policies.

The IKE daemon establishes and installs the following types of SAs:

- A phase 1 SA. For IKEv1, this is known as an ISAKMP SA. For IKEv2, this is known as an IKE SA. Its purpose is to protect communications between IKE peers.
- A phase 2 SA. For IKEv1, this is known as an IPSec SA. For IKEv2, this is known as a child SA. Its purpose is to protect internet protocol (IP) traffic originating from, destined to, or routed by the z/OS TCP/IP stack.

© Copyright IBM Corp. 2000, 2015

The IKE daemon installs three primary types of information in the TCP/IP stack:

Phase 2 SAs
The IKE daemon installs established phase 2 SAs in the TCP/IP stack. On z/OS, the phase 2 SA information that is installed in the TCP/IP stack is referred to as a dynamic tunnel.

Dynamic IP filters
When the IKE daemon installs a dynamic tunnel in the TCP/IP stack, it also installs dynamic IP filters that define what IP traffic can be sent or received through the tunnel. The IKE daemon installs one inbound and one outbound dynamic IP filter with each dynamic tunnel.

Phase 1 SAs
For Sysplex-Wide Security Association (SWSA) support, the IKE daemon also installs phase 1 SA information in the TCP/IP stack. This is only done for SAs established using IKEv1. The IKE daemon only installs phase 1 SAs in a stack that is configured for SWSA support using the DVIPSEC keyword. See z/OS Communications Server: IP Configuration Guide for more information about SWSA support. For information about diagnosing SWSA problems, see "Steps for diagnosing sysplex-wide security association (SWSA) problems" on page 403.

Diagnosing IKE daemon problems

This topic contains information helpful in diagnosing IKE daemon problems.

Initialization problems

When IKE successfully initializes, message EZD1046I is issued. If the IKE daemon fails to initialize, message EZD1045I or EZD1049I is issued. Common initialization problems include:

- The IKE daemon load module is not APF-authorized. The IKE daemon load module must be APF-authorized. The symptom for this problem is the following message: `EZD0986I IKE is not APF authorized`.

  To correct this problem, ensure that the IKE daemon load module is in an APF-authorized library, and then restart the IKE daemon.
- IKE cannot create the /var/ike or /var/sock directories. The IKE daemon attempts to create the /var/ike and /var/sock directories at initialization. If IKE cannot create either of these directories, then initialization fails. If this problem occurs, one of the following messages is issued:

  ```
  EZD1045I IKE initialization error : mkdir /var/ike failed
  EZD1045I IKE initialization error : mkdir /var/sock failed
  ```

  To correct this problem, ensure that the /var directory is mounted as read/write. Additionally ensure that either the IKE daemon has permission to create these directories or that these directories are already created and that the IKE daemon has access to them. If the problem still occurs, contact IBM for additional assistance.
- If the IKE daemon is configured in FIPS 140 mode, ICSF must be active when the IKE daemon is started. If ICSF is not active, message EZD2017I is issued. To correct this problem, start ICSF and then start the IKE daemon. ICSF must be configured in one of several FIPS operational modes when started.

Problems establishing security associations

This topic describes problems in establishing security associations and offers guidance on what steps to take to overcome these problems.

Common problems

Table 17 lists common problems in establishing security associations.

Table 17. Establishing security associations problems

Problem	Symptom	Cause/response
Cannot send or receive packets on UDP ports 500 or 4500	Message EZD1065I was issued. When filter logging is active, message EZD0815I is issued, showing packets to UDP port 500 or UDP port 4500 that were denied.	The IKE daemon communicates using UDP ports 500 and 4500 for IPv4. The IKE daemon communicates using UDP port 500 for IPv6. See "Steps for verifying IP routing to a destination when not using policy-based routing" on page 29 to verify that the IKE daemon is running and bound to ports 500 and 4500. A filter rule must be configured to permit inbound UDP traffic from any source port to destination ports 500 and 4500. A filter rule must be configured to permit outbound UDP traffic from source ports 500 and 4500 to any destination port. Use the **ipsec -f display** command to confirm there is a filter rule installed in the stack that permits receiving traffic from any source UDP port to destination UDP ports 500 and 4500. Also confirm that there is a filter rule that permits receiving traffic from source UDP ports 500 and 4500 to any destination port. Activate filter logging for these rules so that you can observe packets sent on source ports 500 and 4500 and received on destination ports 500 and 4500 in the syslog. For information about the **ipsec** command, see z/OS Communications Server: IP System Administrator's Commands. See z/OS Communications Server: IP Configuration Guide for general information about configuring IP filters.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
Pre-shared key mismatch	Message EZD0965I was issued.	If IKE is using pre-shared key mode authentication and it cannot interpret a decrypted message that it has received, then message EZD0965I is issued, indicating a likely pre-shared key mismatch. In main mode, the responder gets the message upon receipt of message 5. In aggressive mode, the initiator gets the message upon receipt of message 2. EZD0965I can also be issued if IKE receives a corrupted message even though the pre-shared keys match. If the remote peer cannot decrypt the message that was sent by IKE because of a pre-shared key mismatch, the local symptom is that IKE retransmits the first encrypted message of the exchange. Review the pre-shared key configuration on the local and remote system and ensure that the keys match. **Tip:** The keys might be represented differently (for example, ASCII or EBCDIC) on the local and remote system.
Failure accessing local certificate repository	One of the following messages was issued: • EZD0990I • EZD1030I	For the IKE daemon to support RSA signature mode authentication using a local certificate repository, the daemon must be able to access certificates on the SAF key ring. IKE issues message EZD0990I to indicate that RSA signature mode is supported or EZD1030I if RSA signature mode is not supported for a given stack using the key ring. See the messages to determine the appropriate response. The key ring is specified on the KeyRing parameter in the IkeConfig statement. When configuring with the IBM Configuration Assistant for z/OS Communications Server GUI, the key ring is specified on the key ring database field on the IPSec: IKE Daemon Settings panel.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
Failure accessing the network security services (NSS) server	One of the following messages was issued: • EZD1136I • EZD1137I • EZD1138I • EZD1905	For the IKE daemon to support digital signature mode authentication using IPsec certificate services, it must be able to connect to a network security server. IKE issues the following messages: • EZD1136I to indicate that it has connected to a network security server • EZD1137I to indicate that it is not connected to a network security server. • EZD1138I to indicate that it is connecting to a network security server. For the IKE daemon to support digital signature mode authentication for IKEv2 security association activation, IPSec certificate services must be provided by a z/OS V1R12 or later NSS server. If the IKE daemon is not connected to an NSS server, or it is not configured for certificate services, or the NSS is not V1R12 or later, EZD1905 is issued for each IKEv2 SA activation attempt that requires digital signature mode authentication.The network security services server is specified on the NetworkSecurityServer and NetworkSecurityServerBackup parameters on the IkeConfig statement. When configuring with the IBM Configuration Assistant for z/OS Communications Server GUI, the network security server is specified on the server setting in the NSS perspective.
RSA signature authentication failure - missing certificate in the local certificate repository	Message EZD1037I was issued.	Check the syslog to determine whether message EZD1037I was issued. If the IKE daemon cannot locate a certificate that is needed for RSA signature mode authentication, message EZD1037I is issued. • Display the certificates on the SAF key ring. • Ensure that all the certificates on the key ring that are to be used by the IKE daemon include a digital signature. • If you are using RACF, make sure that the trust status of the certificates is TRUST or HIGHTRUST. Use the IKE daemon IkeSyslogLevel 64 to display the contents of the IKE daemon's certificate caches and ensure that the certificates that you want are included in the caches.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
RSA signature authentication failure because of identity mismatch	One of the following messages was issued: • EZD0981I • EZD1075I	Check the syslog to determine whether message EZD0981I or EZD1075I was issued. If the identity that is contained within a received certificate does not match the identity that is configured on the RemoteSecurityEndPoint statement, message EZD0981I is issued. If the peer detects such a mismatch, it might send an "Invalid ID information" notification. If IKE receives such a notification, message EZD1075I is issued. See the messages to determine the appropriate response.
RSA signature authentication failure because of a local certificate verification or authentication failure	One of the following messages was issued: • EZD0902I • EZD0903I	Check the syslog to determine whether message EZD0902I or EZD0903I was issued. If the certificate that is received from a peer cannot be verified, message EZD0902I is issued. If the certificate that is received from the peer cannot be authenticated, message EZD0903I is issued. See the messages to determine the appropriate response. Activate IkeSyslogLevel 64 to get additional diagnostic information that relates to RSA signature mode authentication. The IKE daemon syslog level is set in the IkeSyslogLevel parameter in the IkeConfig statement. When configuring with the IBM Configuration Assistant for z/OS Communications Server GUI, the IKE Daemon Syslog settings are accessed from the IPSec: IKE Daemon Settings panel. IKE maintains a separate cache for Certificate Authority (CA) certificates and security endpoint certificates. When IkeSyslogLevel 64 is active, the contents of the certificate caches are displayed when they are built or rebuilt. See z/OS Communications Server: IP Configuration Reference for information about setting the IkeSyslogLevel, or see the online help in the IBM Configuration Assistant for z/OS Communications Server. **Tip:** The name of the key ring is case-sensitive.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
RSA signature authentication failure - IPsec certificate services failure	Message EZD1139I was issued.	Check the syslog to determine whether message EZD1139I was issued. The EZD1139I message will be issued if the network security server failed to locate a certificate, could not verify a digital signature, or could not create a digital signature for RSA signature mode authentication. See the messages to determine the appropriate response.
Failure to locate phase 1 policy	Message EZD0917I was issued.	In order for IKE to establish a phase 1 SA, it must first locate an applicable phase 1 policy. KeyExchangeRules encapsulate phase 1 policy for IKE. KeyExchangeRules are classified according to a 4-tuple that is comprised of LocalSecurityEndpoint Location, LocalSecurityEndpoint Identity, RemoteSecurityEndpoint Location, and RemoteSecurityEndpoint Identity. When IKE needs to locate a KeyExchangeRule statement, it performs a search of the configured KeyExchangeRules statements, supplying specific values or `Any` for each parameter of the classification 4-tuple. When configuring with the IBM Configuration Assistant for z/OS Communications Server the following are configured in each Connectivity Rule: • Local Security End Point Location • Local Security End Point Identity • Remote Security End Point Location • Remote Security End Point Identity • Key Exchange Settings It is also possible in the GUI to configure a single Local Security End Point Location and Identity for an entire TCP/IP stack. If IKE fails to locate an applicable KeyExchangeRule statement, message EZD0917I is issued that lists the classification 4-tuple. Use the **pasearch -v k -r** command to review the configured KeyExchangeRules statement. If there is no KeyExchangeRule statement that corresponds to the classification 4-tuple that is given on the EZD0917I message, configure a new KeyExchangeRule statement as needed. See the messages for EZD0917I for more information.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
IKE version mismatch	Message EZD1753I was issued.	The IKE daemon sent an IKEv2 initiation request to an IKEv1 peer, and it responded with an IKEv1 message, probably a rejection of the request. The local policy must have HowToInitiate IKEv2, but the peer does not support IKEv2. Change the value on HowToInitiate to Main or Aggressive, to cause the IKE daemon to use IKEv1 flows for tunnel initiation. Alternatively, have the peer IKE node initiate the tunnel. z/OS IKED accepts either IKEv1 or IKEv2 initiation requests, as long as the other attributes being negotiated are properly configured.
Phase 1 policy mismatch	Message EZD1093I or EZD1075I was issued.	The ISAKMP initiator and responder must agree on phase 1 policy to successfully complete negotiation of a phase 1 security association. If the IKE daemon rejects the phase 1 policy that is proposed by an ISAKMP peer, it issues message EZD1021I, which indicates the KeyExchangeRule and KeyExchangeAction statements that were in effect when the mismatch occurred. Message EZD1093I is issued, which indicates why the mismatch occurred. When configuring with the IBM Configuration Assistant for z/OS Communications Server, the Key Exchange Settings are set in each Connectivity Rule. If the IKE daemon proposes phase 1 policy that the ISAKMP peer rejects, the ISAKMP peer should send a notification message to the IKE daemon. If the IKE daemon receives such a notification, it issues message EZD1075I. For more information, see the EZD1075I message documentation in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). If the peer is a z/OS IKE daemon, it issues the EZD1021I and EZD1093I messages as described above. If the peer is not a z/OS IKE daemon, consult the documentation for the ISAKMP peer product to determine why it rejected the proposal. In the case of a mismatch, a No proposal chosen notification is expected from the peer.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
Failure to locate phase 2 policy	Message EZD1795I was issued	In order for IKE to establish a phase 2 SA, it must first locate an applicable phase 2 policy. Phase 2 policy for the IKE daemon is comprised of IpFilterRule and IpDynVpnAction statements. The first step in locating a phase 2 policy for the IKE daemon is to locate an IpFilterRule statement that matches the traffic to be protected and includes a reference to an IpDynVpnAction statement. If IKE cannot find an applicable IpFilterRule statement, message EZD1795I is issued, which indicates the traffic that was to be protected. When configuring with the IBM Configuration Assistant for z/OS Communications Server, the Connectivity Rules are used to locate the phase 2 policies. The Connectivity Rules contain the local and remote data end points, and which type of traffic is protected by Security Levels implementing dynamic tunnels. See the messages to determine the appropriate response. See “Steps for verifying IP security and defensive filter operation” on page 741, supplying the IP traffic characteristics identified on the EZD1795I message.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
Phase 2 policy mismatch	Message EZD1022I, EZD1093I, or EZD1075I was issued.	The ISAKMP initiator and responder must agree on phase 2 policy to successfully complete negotiation of a phase 2 security association. If the IKE daemon rejects the phase 2 policy that is proposed by an ISAKMP peer, it issues message EZD1022I, which indicates the IpFilterRule and IpDynVpnAction statements that were applied. When configuring with the IBM Configuration Assistant for z/OS Communications Server, the Connectivity Rules are used to locate the phase 2 policies. The Connectivity Rules contain the local and remote data end points, and which type of traffic is protected by Security Levels implementing dynamic tunnels. Message EZD1093I is issued indicating why the mismatch occurred. Check the syslog to determine whether message EZD1075I was issued. If the IKE daemon proposes a phase 2 policy that the ISAKMP peer rejects, the ISAKMP peer should send a notification message to the IKE daemon. If the IKE daemon receives such a notification, it issues message EZD1075I. Review the diagnostic data at the ISAKMP peer to determine why the peer rejected the proposal. See the EZD1075I message documentation for more information. In the case of a mismatch, a No proposal chosen notification is expected from the peer.

Table 17. Establishing security associations problems (continued)

Problem	Symptom	Cause/response
AES encryption/decryption failure	Message EZD0918 or EZD1109 was issued.	If IKE is using AES for phase1 or phase2 encryption and decryption, it calls Integrated Cryptographic Service Facility (ICSF) to do the actual cryptography. If ICSF has not been started the cryptography cannot be performed, and IKE cannot encrypt or decrypt messages using AES. This can happen any time during an informational exchange, any time during a phase 2 exchange, in message 5 of main mode if acting as the initiator, or in message 6 of main mode if acting as the responder. The return and reason codes for an ICSF failure are output in message EZD0918 (encryption) or message EZD1109 (decryption). If the return code is C(12) and the reason code is 0, this normally means that ICSF has not been started and therefore cannot perform the necessary cryptography. Ensure that ICSF is started so that the IKE daemon can perform AES cryptography. If the return code is C(12) and the reason code is 8, this normally means that the installed version of ICSF does not support AES. The Security Level Feature of ICSF is required (FMID HCR7706 or higher) for AES support.

Network security services client problems:
IKED can be configured to request network security services (NSS) from an NSS server. The following table lists common problems when IKED, running as an NSS client, is unable to obtain services from the NSS server.

Table 18. Common problems when IKED, running as an NSS client, is unable to obtain services from the NSS server

Problem	Symptom	Cause/response
SSL is not properly configured for IKED, running as an NSS client, to connect to the NSS server.	When AT-TLS is not enabled or is misconfigured on the TCP/IP stack used by IKED or the NSS server, IKED issues message EZD1149I indicating that the connection is not secure.	AT-TLS must be enabled on both the client and server stacks with the TCPCONFIG TTLS statement in the TCP/IP profile. AT-TLS policies must be defined for both the client and the server to secure the connection. See "Define AT-TLS policy to protect communication with an NSS server" in z/OS Communications Server: IP Configuration Guide. If AT-TLS is enabled and the definitions are configured on the client and server stacks but EZD1149I is still displayed then see Chapter 28, “Diagnosing Application Transparent Transport Layer Security (AT-TLS),” on page 707.

Table 18. Common problems when IKED, running as an NSS client, is unable to obtain services from the NSS server (continued)

Problem	Symptom	Cause/response
The userid used for the IKED connection to the NSS server has insufficient authority to connect.	IKED issues message EZD1139I with reason code NSSRsnUserAuthentication. For example: `EZD1139I Request type NSS_ConnectClientReqToSrv with correlator ID 00000000000000040000000000000000 for stack TCPCS2 failed - return code EACCES reason code NSSRsnUserAuthentication`	The IKED connection to the NSS server requires configuration of a valid userid and password or passticket on the NssStackConfig statement in the IKED configuration file.
The userid used for the IKED connection to the NSS server has insufficient authority to access services requested.	IKED issues messages indicating which requested services are not available. For example: • EZD1145I The network security certificate service is not available for stack TCPCS2 • EZD1147I The network security remote management service is not available for stack TCPCS2	The following SAF resource permissions are required to access network security services: • EZB.NSS.sysname.clientname.IPSEC.CERT • EZB.NSS.sysname.clientname.IPSEC.NETMGMT These resources must be defined on the NSS server system and the userid configured on the NssStackConfig statement in the IKED configuration file must be permitted read access to them.
IKED fails to retrieve certificates from the NSS server.	IKED syslog daemon traces may show that no cache entries were received from the NSS server. For example: `IKE: Initializing CA Cache with 0 entries for stack TCPCS2` Dynamic tunnel negotiations using RSA signature mode fail.	The following SAF resource permissions are required to access certificates from the NSS server: • EZB.NSSCERT.sysname.mappedlabelname.CERTAUTH • EZB.NSSCERT.sysname.mappedlabelname.HOST These resources must be defined on the NSS server system and the userid configured on the NssStackConfig statement in the IKED configuration file must be permitted read access to it. See "Steps for authorizing resources for NSS" in z/OS Communications Server: IP Configuration Guide.
IKED does not attempt to connect to the NSS server for a given stack.	IKED does not issue message EZD1138I for the given stack.	A valid NssStackConfig statement is required for each stack to use NSS. See IKE daemon in z/OS Communications Server: IP Configuration Reference for information about configuring the NssStackConfig statement.
IKED connects to NSSD but cannot use NSS IPSec certificate services.	Message EZD1916I was issued.	IKED is configured in FIPS 140 mode, but the NSS server is not. Therefore, IKED cannot use the NSS certificate services provided by the NSS server because the cryptographic operations performed by the NSS server on behalf of IKED will not be performed in a manner consistent with FIPS 140 requirements. IKED remains connected to the NSS server so it can use the NSS remote management services.

NAT traversal considerations

- NAT traversal support must be enabled.

  By default, NAT traversal support on z/OS is disabled. To enable NAT traversal support do one of following:

- Specify a value of `Yes` for the AllowNat parameter of the KeyExchangeAction statement used when negotiating with a remote security endpoint that you want to perform NAT traversal with.
- Specify a value of `Yes` for the AllowNat parameter of the KeyExchangePolicy statement. Verify that the AllowNat parameter is not specified as `No` on the KeyExchangeAction statement used when negotiating with the remote security endpoint that you want to perform NAT traversal with.

Use care when using the latter method. The AllowNat parameter specified on the KeyExchangePolicy statement becomes the default AllowNat setting for KeyExchangeAction statements that do not specify the AllowNat parameter. See z/OS Communications Server: IP Configuration Reference for more details concerning the AllowNat parameter. When configuring with the IBM Configuration Assistant for z/OS Communications Server, you configure whether to allow NAT traversal processing on the Stack Level Settings panel. This setting can be overridden in each Connectivity Rule.

The AllowNat field contained in the output of the **ipsec -k display** command can be used to determine whether NAT Traversal support was enabled for a phase 1 negotiation.

Changes made to the AllowNat parameter do not impact existing phase 1 security associations. Existing phase 1 security associations must be refreshed before any changes to the AllowNat setting are honored. There are no configuration options to enable or disable NAT traversal for an IPSec security association. The state of NAT traversal for an IPSec security association is determined by the phase 1 security association used when negotiating the IPSec security association.

- The remote security endpoint must support an acceptable version of NAT traversal. z/OS provides limited support for the following levels of NAT Traversal:
  - draft-ietf-ipsec-nat-t-ike-02
  - draft-ietf-ipsec-nat-t-ike-03
  - RFC 3947
  - RFC 3947 with z/OS-only extensions
  - RFC 5996
  - RFC 5996 with z/OS-only extensions

  The remote security endpoint must support one of these levels of NAT traversal.

  You can use the NATTSupportLevel field contained in the output of the **ipsec -k display** command to determine what level of NAT traversal support was used during a phase 1 negotiation. If the NATTSupportLevel is `None`, verify that NAT traversal support is enabled when negotiating with this remote security endpoint. If NAT traversal support was enabled, then the remote security endpoint does not provide an acceptable level of NAT traversal support.
- z/OS does not support NAT traversal for IPv6 traffic.
- z/OS cannot act as a gateway when traversing a NAT.

  The z/OS IKE daemon does not support acting in the gateway role when traversing a NAT. The z/OS IKE daemon is acting as a gateway when the local data endpoint of an IPSec security association is not the same as the IP address used as the local IP address of the protecting phase 1 security association. Message EZD1089I is issued when z/OS is acting as a gateway while traversing a NAT.

  When a NAT is detected between z/OS IKE and a remote security endpoint, all IPSec security associations negotiated with that remote security endpoint must

end in the local z/OS box. Specifically, the local data endpoint of any IPSec security association negotiated when traversing a NAT must be the IP address used as the local IP address of the protecting phase 1 security association. If z/OS is behind a NAT this could be its private address or the public IP address provided by the NAT.

You can use the LocalEndpoint field contained in the **ipsec -k display** command used to determine the local private IP address of the protecting phase 1 security association. The local public IP address of the protecting phase 1 security association is assigned by the NAT box in front of z/OS.

You can use the NATInFrntLclScEndPnt and NATInFrntRmtScEndPnt fields contained in the output of the **ipsec -k display** command to determine whether a NAT was detected between the IKE daemon and the remote security gateway.

- z/OS cannot act as an initiator to a security gateway.

  The z/OS IKE daemon cannot act as the initiator of a phase 2 negotiation for a new IPSec security association when traversing a NAT and the remote security endpoint is acting as security gateway. Messages EZD1090I or EZD1057I are issued in this case. The z/OS IKE daemon can act as the initiator of subsequent refreshes of an existing IPSec security association with a remote security endpoint that is acting as a security gateway.

  A new IPSec security association is the first IPSec security association negotiated for a particular traffic pattern. A remote security endpoint is acting as a security gateway when the remote endpoint of the IPSec security association is not the same as the IP address used as the remote IP address of the protecting phase 1 security association.

- z/OS cannot act as an initiator to a remote security endpoint located behind a NAT device performing network address port translation (NAPT).

  The z/OS IKE daemon does not support initiating the first security association to a remote security endpoint when a NAT has translated the remote security endpoint's port (that is, when IKE detects the existence of a NAPT in front of the remote security endpoint). If this condition is detected during a phase 1 or phase 2 negotiation, the negotiation is terminated.

- z/OS uses only IPv4 identities during IKEv1 phase 2

  During an IKEv1 phase 2 negotiation for a new IPSec security association, the z/OS IKE daemon uses IPv4 ID types to identify the traffic pattern to be protected by the new IPSec security association. When traversing a NAT, other IKE implementations might require the traffic pattern to be specified using a non-IPv4 ID type. The z/OS IKE daemon is not able to act as the initiator of an IKEv1 phase 2 negotiation with such an implementation when creating a new IPSec security association. The z/OS IKE daemon can act as the initiator of subsequent refreshes of an existing IPSec security association with a remote security endpoint utilizing such an IKE implementation.

  When the z/OS IKE daemon acts as the initiator of an IKEv1 phase 2 negotiation to create a new IPSec security association and the remote security endpoint is using an implementation that requires a non-IPv4 ID type, the remote security endpoint rejects the proposal. Some implementations might send an informational notification in this case. The informational notification indicates that the proposal was rejected and why. If an informational notification is received, the z/OS IKE daemon issues message EZD1075I.

- Interoperability considerations when z/OS initiates a phase 2 negotiation to a non-z/OS peer for a host-to-host tunnel that traverses a NAT
  - Host-to-host dynamic tunnel protecting all ports and all protocols, tunnel mode

When z/OS initiates phase 2 negotiation for a new host-to-host dynamic tunnel that protects all ports and all protocols, z/OS uses a traffic pattern consisting of the IP addresses of the local and remote security endpoints. IKE and its peer view the traffic pattern differently. z/OS views the traffic pattern as its private IP address and the peer's public IP address. The peer views the traffic pattern as the public address of z/OS and the peer's private address.

A host-to-host dynamic tunnel uses either the transport or tunnel mode of encapsulation. When z/OS initiates a phase 2 negotiation to a non-z/OS peer for a new host-to-host dynamic tunnel that protects all ports and all protocols it is possible that the negotiation succeeds, but it produces an SA that cannot be used to send traffic. This is partially because data protected using tunnel mode SAs have two IP headers. Because both peers have a different view of the traffic pattern, they might not agree on the contents of the inner-most IP header. When both the local and remote peers are z/OS, the SA negotiation should be successful and produce an SA that can be used to send traffic.

– Host-to-host dynamic tunnel protecting specific ports or protocols

  RFC 3947 does not discuss how traffic patterns should be defined when one or more NATs are being traversed. When z/OS initiates a phase 2 negotiation for a new host-to-host dynamic tunnel that protects a specific ports or protocol, it defines the traffic pattern using z/OS private addresses as the local endpoint, if z/OS is behind a NAT, and the peer's public address as the remote endpoint. In this case, the negotiation might fail with a non-z/OS peer, depending on the NAT traversal support of the non-z/OS peer. The negotiation should be successful with a z/OS peer.

  To help identify configurations where there are potential interoperability concerns, three informational messages have been defined. When z/OS initiates a phase 2 negotiation for a UDP encapsulated tunnel mode SA with a non-z/OS peer, message EZD1104I or EZD1105I is issued for IKEv1 and message EZD1924I or EZD1925I is issued for IKEv2. When z/OS initiates a phase 2 negotiation for a UDP-encapsulated tunnel or transport-mode SA for a specific port, protocol, or both, message EZD1107I is issued for IKEv1. In all cases, the negotiation continues.

• SWSA implications

  During VIPA takeover or giveback processing, the IKE daemon attempts to create security associations that existed on the stack that owned the security association prior to the takeover or giveback. These security associations appear as new security associations on the new owning stack.

  The z/OS IKE daemon cannot initiate the creation of new IPSec security associations when the peer is acting as a gateway, or when the peer is behind a NAPT, or when the peer expects a non-IPv4 identity during a quick mode exchange; however, it can act as a responder in these cases. VIPA takeover and giveback of such phase 2 security associations are not supported.

  There are also cases when the results might be unpredictable when the z/OS IKE daemon initiates a new host-to-host SA negotiation to a non-z/OS peer. These cases include:

  – z/OS IKE initiates a new host-to-host UDP encapsulated tunnel mode SA to a non-z/OS peer
  – z/OS IKE initiates a new host-to-host UDP encapsulated SA for a specific port, protocol, or both to a non-z/OS peer

  It is expected that IKE can always act as a responder in these cases. If such SAs exist when a VIPA takeover or giveback occurs, the IKE daemon attempts to re-establish these security associations. The results of these attempts are unpredictable. This can result in a disruption of traffic until new SAs are created

by the remote security endpoint. The IKE daemon still sends delete notifications informing the remote security endpoint that the security associations are no longer valid.

- Remapping of a remote security endpoint's address

  When a remote security endpoint is behind a NAT, the NAT maps the private IP address of the remote security endpoint to a public IP address. This mapping could expire as a result of inactivity or a new mapping could be created due to a reboot of the NAT device. In such cases, the public IP address of the remote security endpoint might change.

  In the cases where NAT performs port translation (NAPT), the IP address or port or both might change.

  If the IKE daemon or stack detects such a change while there are one or more security associations (SAs) with that remote security endpoint, the IKE daemon attempts to verify the new IP address and port pair. It does this by initiating the creation of a new phase 1 SA using the remote security endpoint's new IP address and port pair. Message EZD1086I is issued when this negotiation starts. If this negotiation is successful, the IKE daemon issues message EZD1087I and all phase 1 and IPSec security associations with that remote security endpoint using the old address are deleted. IKEv2, as specified in RFC 5996, requires an attempt to create a phase 2 SA when creating a phase 1 SA. When the IKE daemon attempts to verify a new IP address and port pair with IKEv2, one or more phase 2 SAs can be initiated when the phase 1 SA is initiated. The number of phase 2 SA initiations attempted is determined by the number of phase 2 SAs belonging to the phase 1 SA that corresponds to the remote security endpoint for which an IP address or port change was detected.

- NAT keepalive timer

  When a z/OS is behind a NAT, the NAT maps its private IP addresses to public IP addresses. A static NAT mapping does not expire. A dynamic NAT mapping can expire as a result of inactivity. To prevent the expiration of this mapping, the stack occasionally sends messages known as NAT keepalive messages. If these messages are not sent frequently enough, the NAT device could expire the mapping of any z/OS private IP addresses to public IP addresses. Such a remapping could be disruptive to existing IPSec traffic.

  The frequency of message transmission is defined by the NatKeepAliveInterval value on the KeyExchangePolicy statement. See z/OS Communications Server: IP Configuration Reference for more details. When configuring with the IBM Configuration Assistant for z/OS Communications Server, the NAT keep alive interval is specified on the Stack Level Settings panel.

  A NAT keep alive is a 1–byte UDP message sent to a remote security endpoint using the UDP encapsulation ports. The sent byte is set to x'FF'. Figure 29 shows a NAT keep alive message:

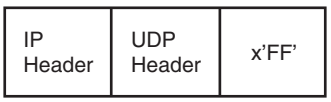

Figure 29. NAT keep alive message

- Multiple remote security endpoints sharing the same phase 1 identity

  During a phase 1 negotiation the remote security endpoint sends its identity in an ID payload. The IKE daemon can manage multiple remote security endpoints using the same ID when those endpoints are not behind a NAT. However, when

a remote security endpoint is behind a NAT it must use a unique phase 1 identity. If a second remote security endpoint behind a NAT attempts to use a phase 1 identity already in use by another remote security endpoint behind a NAT, the IKE daemon detects this as a remapping of the first remote security endpoint's IP address.

When multiple remote security endpoints behind a NAT share the same phase 1 identity, messages EZD1086I and EZD1087I might be repeatedly issued.

- Responding to phase 1 main mode SA negotiations with multiple remote security endpoints behind a NAPT

  When acting as a responder in main mode SA negotiations, the z/OS IKE daemon must agree to key exchange parameters before the remote security endpoint identity is known. The key exchange policy is searched to match on a KeyExchangeRule based upon the IP addresses of the local and remote security endpoints. Different remote security endpoints located behind an NAPT might use the same public IP address. This can cause a policy mismatch if the KeyExchangeRule settings for those remote security endpoints do not match. A policy mismatch will cause EZD1093I to be issued to syslog.

  To prevent this situation, do the following:

  Configure a single KeyExchangeRule to represent all the remote security endpoints behind the NAPT device. The remote security endpoints, represented by the same public address, must use the same security parameters. The remote security endpoints' policy must be configured with the same security parameters as well.

Resource constraint problems

This topic describes problems that can occur when the IKE daemon experiences shortages of certain system resources. The following table lists system resources that can become constrained when large numbers of remote IKE peers simultaneously attempt to negotiate IPSec tunnels with z/OS. For each resource, instructions are also provided for alleviating the constraints.

Table 19. Resource constraint problems

Problem	Symptom	Cause/response
Maximum number of messages on system message queue exceeded	Message EZD0970I was issued, which indicates that the C/C++ runtime library function msgsnd() failed with the following error information: 021D - 112 | 070B031C | EDC5112I Resource temporarily unavailable.	The IKE daemon uses system message queues to internally route work between different threads. In this case, one of these queues reached its maximum capacity as defined by the IPCMSGQMNUM parameter of the SETOMVS command. This situation occurs when thousands of IKE peers simultaneously attempt to negotiate new or refresh existing security associations with z/OS. To resolve the problem, use the SETOMVS command with the IPCMSGQMNUM parameter to specify a higher message queue capacity.

Table 19. Resource constraint problems (continued)

Problem	Symptom	Cause/response
IKED memory limits exceeded	Message EZD0963I was issued multiple times, which indicates that the IKE daemon cannot allocate memory.	This situation occurs when thousands of IKE peers simultaneously attempt to negotiate new or refresh existing security associations with z/OS. To resolve the problem, increase the amount of virtual storage available to the IKE daemon. For instructions on how to ensure that the IKE daemon's user ID has enough system resources, see Step 2: Ensure the IKE daemon's user ID has enough system resources in z/OS Communications Server: IP Configuration Guide.
UDP queue limits for IKE port exceeded	• An unusually high percentage of receive errors is displayed in the UDP statistics section of the NETSTAT STATS/-S report. • In some cases when no UDPQUEUELIMIT is specified in the TCPIP profile, you might also see message IST2273E.	The IKE daemon communicates by using UDP ports 500 and 4500. If UDP traffic to one of these ports overruns the capacity of the UDP queue in the TCP/IP stack, you might encounter one or both of the listed symptoms. The capacity of the UDP queue is defined by the UDPQUEUELIMIT parameter of the UDPCONFIG statement in the TCPIP profile data set. When this symptom happens, the TCP/IP stack discards inbound UDP packets in a randomized way until the constrained queue condition is relieved. This situation occurs when thousands of IKE peers simultaneously attempt to negotiate new or refresh existing security associations with z/OS. To resolve the problem, enable Intrusion Detection Services (IDS) for the stack and define a UDP Traffic Regulation rule for the IKE daemon's UDP ports with the VERY_LONG queue size. For more information about IDS and UDP Traffic Regulation, see Traffic regulation policies in z/OS Communications Server: IP Configuration Guide.

Abends

Messages and error-related information should be sent to the system console when an abend occurs during IKE daemon processing. A dump of the error is needed unless the symptoms match a known problem. System dumps of IKE include Language Environment data. The Language Environment IPCS verbexit LEDATA can be used to format this information. See z/OS Language Environment Debugging Guide for more information. The following command is a sample IPCS**verbexit ledata** command:

```
verbx ledata 'asid(68) tcb(007E5E88) ceedump nthreads(*)'
```

Tip: In this example, the IKE asid is 0x68 and the address of the abended IKE TCB is 0x007E5E88.

IKE daemon debug information

Additional IKE daemon debug information can be sent to the syslog by using the IkeSyslogLevel and PagentSyslogLevel parameters in the IKE configuration file.

Obtaining syslog debug information for the IKE daemon

The IkeSyslogLevel parameter in the IKE configuration file controls the level of IKE internal debug information that is sent to syslog. When configuring with the IBM Configuration Assistant for z/OS Communications Server, use the IKE Daemon Settings panel to configure the level of IKE internal debug information that is sent to syslog.

The IKE syslog level value should be set above 1 only when diagnosing a problem; levels above 1 impact IKE performance. Level 8 and Level 16 have the greatest performance impact because they affect processing on each UDP datagram IKE sends and receives.

IKE Syslog level values can be combined. See z/OS Communications Server: IP Configuration Reference or the IBM Configuration Assistant for z/OS Communications Server's online help for more information.

Note: When IKED logs messages through syslogd, the messages are immediately preceded by the parenthesized thread identifier of the IKED thread from which they were issued. For example, the following message was logged by thread 1:

```
Jul 28 15:10:47 mvs046 IKE: (00000001) EZD1061I IKE connecting to PAGENT
```

Obtaining debug information using PagentSyslogLevel

IKE uses the Policy API (PAPI) to communicate with the Policy Agent and manipulate policy information it has obtained from the Policy Agent. The PagentSyslogLevel parameter in the IKE configuration file controls the level of debug information that is sent to syslog when IKE uses PAPI. When configuring with the IBM Configuration Assistant for z/OS Communications Server, the Policy Agent Syslog events are configured from the IKE Daemon Settings panel. The Policy Agent Syslog level value should be set above 0 only at the direction of IBM service as it impacts IKE performance. For more information about setting the Policy Agent Syslog levels, see z/OS Communications Server: IP Configuration Reference or the IBM Configuration Assistant for z/OS Communications Server online help.

TCP/IP services component trace for the IKE daemon

z/OS CS provides component trace support for the IKE daemon. This section describes how to specify IKE daemon trace and formatting options. For short descriptions of other tracing procedures, such as displaying trace status, see Chapter 5, “TCP/IP services traces and IPCS support,” on page 45. For detailed information, see the following information:

- z/OS MVS Diagnosis: Tools and Service Aids for information about component trace procedures.
- z/OS MVS Initialization and Tuning Reference for information about the component trace SYS1.PARMLIB member.
- z/OS MVS System Commands for information about commands.
- z/OS MVS Programming: Authorized Assembler Services Guide for procedures and return codes for component trace macros.

Using CTRACE

You can specify component trace options at TCP/IP initialization or after TCP/IP has initialized.

Table 20 lists the IKE daemon trace options.

Table 20. IKE daemon trace options

Trace Event	Description
ALL	Trace all types of records. This option slows performance.
MINIMUM	Trace the IKE daemon's minimum level of tracing.
INIT	Trace IKE daemon initialization information.
TERM	Trace IKE daemon termination information.
EXCEPT	Trace IKE daemon exception information.
CONFIG	Trace IKE daemon configuration information.
WORKUNIT	Trace IKE workunit information.
SERIAL	Trace IKE serialization information.
IKE	Trace IKE protocol information.
CRYPTO	Trace IKE cryptographic information.
OPMSGS	Trace IKE operator messages.
LOGMSGS	Trace IKE syslog messages.
MSGQ	Trace IKE message queue information.
TIMER	Trace IKE timer information.
SOCKETS	Trace IKE socket information.
IOCTL	Trace IKE IOCTL call information.
REQUESTS	Trace IKE request information.
FLOW	Trace IKE code flow information.
STORAGE	Trace IKE storage information.
EVENT	Trace IKE event information.
POLICY	Trace IKE policy information.
CONTROL	Trace IKE daemon control information.
MISC	Trace IKE miscellaneous information.
DEBUG	Trace IKE debugging information.

Enabling CTRACE at IKE daemon startup

A default minimum component trace is always started during IKE daemon initialization. Use a parmlib member to customize the parameters that are used to initialize the trace. The default IKE daemon component trace parmlib member is the SYS1.PARMLIB member CTIIKE00. The parmlib member name can be changed using the IKED_CTRACE_MEMBER environment variable.

Tip: The IKE daemon reads the IKED_CTRACE_MEMBER environment variable only during initialization. Changes to IKED_CTRACE_MEMBER after daemon initialization have no affect.

For a description of trace options, see Table 20.

Restriction: In addition to specifying the trace options, you can also change the IKE daemon trace buffer size. The buffer size can be changed only at IKE initialization and has a maximum of 256 MB.

If the CTIIKE00 member or the member that is specified in IKE_CTRACE_MEMBER is not found when starting the IKE daemon, the following message is issued:

```
IEE5381 memberName MEMBER NOT FOUND IN PARMLIB
```

When this occurs, the IKE daemon component trace is started with a buffer size of 1 MB and the MINIMUM tracing option.

```
/*********************************************************************/
/*                                                                   */
/* z/OS Communications Server                                        */
/* SMP/E Distribution Name: CTIIKE00                                 */
/*                                                                   */
/* PART Name: CTIIKE00                                               */
/*                                                                   */
/*                                                                   */
/* Copyright:                                                        */
/* Licensed Materials - Property of IBM                              */
/* 5694-A01                                                          */
/* (C) Copyright IBM Corp. 2005                                      */
/* Status:  CSV1R7                                                   */
/*                                                                   */
/*                                                                   */
/* DESCRIPTION = This parmlib member causes component trace for      */
/* the TCP/IP IKE application to be initialized                       */
/* with a trace buffer size of 1M                                    */
/*                                                                   */
/* This parmlib members only lists those TRACEOPTS                   */
/* values specific to IKE. For a complete list                       */
/* of TRACEOPTS keywords and their values see                        */
/* z/OS MVS Initialization and Tuning Reference.                     */
/*                                                                   */
/*                                                                   */
/*                                                                   */
/*********************************************************************/
 TRACEOPTS
 /* ---------------------------------------------------------------- */
 /* Optionally start external writer in this file (use both          */
 /* WTRSTART and WTR with same wtr_procedure)                        */
 /* ---------------------------------------------------------------- */
 /* WTRSTART(wtr_procedure)                                          */
 /* ---------------------------------------------------------------- */
 /* ON OR OFF: PICK 1                                                */
 /* ---------------------------------------------------------------- */
 ON
 /* OFF                                                              */
 /* ---------------------------------------------------------------- */
 /* BUFSIZE: A VALUE IN RANGE 128K TO 256M                           */
 /* CTRACE buffers reside in IKE daemon Private storage              */
 /* which is in the regions address space.                           */
 /* ---------------------------------------------------------------- */
 BUFSIZE(1M)
 /* WTR(wtr_procedure)                                               */
 /* ---------------------------------------------------------------- */
 /* OPTIONS: NAMES OF FUNCTIONS TO BE TRACED, OR "ALL"               */
 /* ---------------------------------------------------------------- */
 /* OPTIONS(          */
 /*   'ALL     '      */
 /* , 'MINIMUM '      */
 /* , 'INIT    '      */
```

```
/* , 'TERM    '      */
/* , 'EXCEPT  '      */
/* , 'CONFIG  '      */
/* , 'WORKUNIT'      */
/* , 'SERIAL  '      */
/* , 'IKE     '      */
/* , 'CRYPTO  '      */
/* , 'OPMSGS  '      */
/* , 'LOGMSGS '      */
/* , 'MSGQ    '      */
/* , 'TIMER   '      */
/* , 'SOCKETS '      */
/* , 'IOCTL   '      */
/* , 'REQUESTS'      */
/* , 'FLOW    '      */
/* , 'STORAGE '      */
/* , 'EVENT   '      */
/* , 'POLICY  '      */
/* , 'CONTROL '      */
/* , 'MISC    '      */
/* , 'DEBUG   '      */
/* )                 */
```

Figure 30. SYS1.PARMLIB member CTIIKE00

Steps for enabling the CTRACE at IKE daemon startup

Enable the CTRACE at IKE daemon startup.

Procedure

Perform the following steps:

1. Edit the CTIIKE00 parmlib member and specify TRACEOPTS ON, the buffer size you want with the BUFSIZE() parameter and the CTRACE options you want. To direct the CTRACE to an external writer, also specify the name of the writer JCL procedure in the WTR() parameter. See the example CTIIKE00 parmlib member.
2. Start the IKE daemon.

Steps for disabling the CTRACE at IKE daemon startup

Disable the CTRACE at IKE daemon startup.

Procedure

Perform the following steps:

1. To disable the CTRACE at IKE daemon startup, edit the CTIIKE00 parmlib member and specify TRACEOPTS OFF.
2. Start the IKE daemon.

Step for enabling the CTRACE after the IKE daemon is started

After the IKE daemon is started, enable the CTRACE to an internal buffer or to an external writer.

Procedure

Perform the appropriate step:

- Issue the following console commands to enable the CTRACE to an internal buffer:

```
TRACE CT,ON,COMP=SYSTCPIK,SUB=(iked_jobname)
R xx,OPTIONS=(option[,option2...]),END
```

- Issue the following console commands to enable the CTRACE to an external writer:

```
TRACE CT,WTRSTART=writer_proc
TRACE CT,ON,COMP=SYSTCPIK,SUB=(iked_jobname)
R xx,OPTIONS=(option[,option2...]),WTR=writer_proc,END
```

Step for disabling the CTRACE after the IKE daemon is started

After the IKE daemon is started, disable the CTRACE to an internal buffer or to an external writer.

Procedure

Perform the appropriate step:

- Issue the following console commands to disable the CTRACE to an internal buffer:

```
TRACE CT,OFF,COMP=SYSTCPIK,SUB=(iked_jobname)
```

- Issue the following console commands to disable a CTRACE to an external writer:

```
TRACE CT,OFF,COMP=SYSTCPIK,SUB=(iked_jobname)
TRACE CT,WTRSTOP=writer_proc
```

Step for displaying the CTRACE status

Display the CTRACE status at the console.

Procedure

Perform the following step:

To display the CTRACE status, issue the following console command:

```
D TRACE,COMP=SYSTCPIK,SUB=(iked_jobname)
```

Enabling CTRACE after IKE daemon initialization

After IKE daemon initialization, you must use the TRACE CT command to change the component trace options. Each time a new Component Trace is initiated, all prior trace options are turned OFF and the new options are put into effect. You can specify the trace options with or without the parmlib member. See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45.

Formatting IKE daemon trace records

You can format component trace records using IPCS panels or a combination of the IPCS panels and the CTRACE command, either from a dump or from external-writer files. See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45 for details.

Enter any combination of the following values as options to filter the CTRACE entries. The options must be entered using the following format:

```
TYPE(option[,option]...)
```

You can use any of the options listed in Table 20 on page 352, except ALL and MINIMUM.

Chapter 10. Diagnosing network security services (NSS) server problems

This topic describes how to diagnose network security services (NSS) server problems, and contains the following information:

- "Overview of diagnosing NSS server problems"
- "Network security services server debug information" on page 367
- "TCP/IP services component trace for the network security services (NSS) server" on page 377
- "Steps for enabling the CTRACE at network security service (NSS) server startup" on page 379

Overview of diagnosing NSS server problems

The NSS server provides network security services for one or more network security enforcement points. A component that requests network security services from the network security services server is called a network security client or NSS client. Problems with the network security services server may be categorized as follows:

- Network security services server configuration problems
- Network security services server internal problems
- Network security services server problems interacting with an external component such as a network security client or the Secure Access Facility (SAF).

The NSS server provides log output using syslogd and internal trace information using component trace (CTRACE). The log output is sufficient for diagnosing most network security services server problems and is the first place to look if you suspect a problem.

Common NSS server initialization problems

Table 21. Common NSS server initialization problems

Problem	Symptom	Cause/response
The NSS load module is not APF-authorized.	The NSS load module abends. The following message will be logged to the console: `IEF450I NSSD STEP1 - ABEND=S000 U4087 REASON=00000000`	The NSS load module must be APF-authorized.
The NSS socket directory does not exist or else it cannot be created by the NSS server.	When NSS server syslog level 2 is set (NSS_SYSLOG_LEVEL_VERBOSE), debug message DBG0040I is generated. For example: `DBG0040I NSS_VERBOSE Cannot create socket directory /var/sock - rc -1 errno 135 EDC5135I Not a directory. The NSS server will immediately shutdown.`	1. The /var directory must already exist. 2. The /var/sock subdirectory must already exist, or else the userid that the NSS server is running under must have authority to create the /var/sock subdirectory.

© Copyright IBM Corp. 2000, 2015

Table 21. Common NSS server initialization problems (continued)

Problem	Symptom	Cause/response
The NSS server is configured in FIPS 140 mode and ICSF is not active.	Message EZD1390I is issued, and the NSS server stops.	If the NSS server is configured in FIPS 140 mode, ICSF must be active when the NSS server is started. To correct this problem, start ICSF and then start the NSS server. ICSF must be configured in one of several FIPS operational modes when started.

NSS client connection problems

The following table lists common problems when a network security services (NSS) client is unable to obtain services from the NSS server.

Table 22. Common NSS client connection problems

Problem	Symptom	Cause/response
SSL is not properly configured for the NSS client connection to the NSS server. NSS client fails to connect.	When NSS server syslog level 8 is set (NSS_SYSLOG_LEVEL_CLIENTLIFECYCLE), debug message DBG0104I is generated: `DBG0104I NSS_LIFECYCLE NSS connID 1 - the connection is not secure - the connection will be closed`	For NSS IPSec client connections: • AT-TLS must be enabled on both the client and server stacks with the TCPCONFIG TTLS statement in the TCP/IP profile. • AT-TLS policies must be defined to secure the connection. • See "AT-TLS policy" in z/OS Communications Server: IP Configuration Guide. • If AT-TLS is enabled and the definitions are configured on the client and server stacks but DBG0104I is still displayed then see Chapter 28, "Diagnosing Application Transparent Transport Layer Security (AT-TLS)," on page 707. For NSS XMLAppliance client connections: • The NSS XMLAppliance client must support an SSL/TLS negotiation protocol that is compatible with that which is configured for the NSS server. • The NSS server stack must have AT-TLS enabled with the TCPCONFIG TTLS statement in the TCP/IP profile. • AT-TLS policies must be defined for the NSS server stack to secure the connection. See AT-TLS policy in z/OS Communications Server: IP Configuration Guide. • If AT-TLS is enabled on the server stack, and the definition is configured on the server stack but DBG0104I is still displayed then see Chapter 28, "Diagnosing Application Transparent Transport Layer Security (AT-TLS)," on page 707. Configuration of the client's TLS settings are left up to the client application's implementation.

Table 22. Common NSS client connection problems (continued)

Problem	Symptom	Cause/response
The userid used for the NSS client connection to the NSS server has insufficient authority to access services requested.	When NSS server syslog level 2 is set (NSS_SYSLOG_LEVEL_VERBOSE), debug message DBG0032I is generated. For example: `DBG0032I NSS_VERBOSE ServauthCheck (USER2 , EZB.NSS.MVS093.CLIENT2.IPSEC.CERT) rc 4 (DENY) racfRC 4 racfRsn 0`	SAF resource permissions are required to access NSS IPSec services: • EZB.NSS.*sysname.clientname*.IPSEC.CERT • EZB.NSS.*sysname.clientname*.IPSEC.NETMGMT SAF resource permissions are required to access the NSS XMLAppliance services: • `EZB.NSS.sysname.clientname. XMLAPPLIANCE.SAFACCESS` • `EZB.NSS.sysname.clientname. XMLAPPLIANCE.CERT` • `EZB.NSS.sysname.clientname. XMLAPPLIANCE.PRIVKEY` These resources must be defined on the NSS server system and the client userid must be permitted read access to them.
An NSS client appears to be connected to two instances of the NSS server.	For an NSS IPSec client, the ipsec -x display for both NSS servers shows the same client connected. For an NSS client, the nssctl -d for both NSS servers shows the same client connected.	Under normal termination, an NSS client will issue a disconnect to close its connection with the NSS server. In some rare recovery situations, the NSS server may not be aware that a connection with an NSS client has ended. When the client restarts or attempts to reconnect, it is possible it may connect to a different NSS server instance, such as the backup server or an NSS server on another system when the client is connecting on a distributed dynamic VIPA. Use the ipsec -w display on the system running the affected NSS IPSec client to determine which NSS server the IPSec client is actually connected. Optionally, use the Netstat DRop/-D command to close out the old connection on the other NSS server.
NSS clients are failing to connect to the NSS server.	The NSS server issues the EZD1371I console message to indicate the disabled discipline and closes the connection.	The NSS server has been configured to disable the specified discipline. Modify the NSS server configuration to enable the specified discipline. See z/OS Communications Server: IP Configuration Reference for more information about the NSS server configuration.

The following table lists common problems when requests from a network security services (NSS) client fails.

Table 23. Common NSS client request failures

Problem	Symptom	Cause/response
The userid used for the NSS client connection has insufficient authority to access client certificates.	When NSS server syslog level 4 is set (NSS_SYSLOG_LEVEL_CERTINFO), debug message DBG0004I is generated: `DBG0004I NSS_CERTINFO Client MVS093_TCPCS3 connected as userid USER1 is not authorized to profile EZB.NSSCERT.VIC012.NSCLIENT3.HOST associated with matching certificate ( NSCLIENT3 ) for request 00000000000000150000000000000000`	SAF resource permissions are required to access certificates from the NSS server: • `EZB.NSSCERT.`*`sysname.mappedlabelname`*`.HOST` • `EZB.NSSCERT.`*`sysname.mappedlabelname`*`.CERTAUTH` These resources must be defined on the NSS server system and the client userid must be permitted read access to them.
The userid used for the NSS client connection has insufficient authority to access the private keys associated with client certificates.	When NSS server syslog level 4 is set (NSS_SYSLOG_LEVEL_CERTINFO), debug message DBG0004I is generated: `Jun 25 14:54:43 MVS093 NSSD: DBG0004I NSS_CERTINFO Client XML_ClientB8 connected as userid USER198 is not authorized to profile EZB.NSSCERT.MVS093. KEY1024ICSF.PRIVKEY associated with matching certificate ( Key1024ICSF ) for req uest 9987A24B879696844B824BF0F0F10000`	SAF resource permissions are required to access private keys associated with certificates from the NSS server: `EZB.NSSCERT.`*`sysname.mappedlabelname`*`. PRIVKEY` These resources must be defined on the NSS server system and the client userid must be permitted read access to them.

NSS XMLAppliance client API return codes and reason codes

The following table lists and describes the possible return codes and reason codes returned by the NSS server to NSS XMLAppliance clients.

Table 24. NSS XMLAppliance client API return codes and reason codes

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
0	0	No error
EINVAL(121)	NMsRsnBadIdent (1)	Invalid message or record identifier supplied in message. **System Action:** Request is failed but connection remains open. **Response:** Re-issue the request and send a correctly formatted message
EINVAL(121)	NMsRsnBadVersion (2)	Invalid version supplied in message header. **System Action:** Request fails but connection remains open. **Response:** Send a correctly formatted message.
EINVAL(121)	NMsRsnBadType (3)	Unsupported or unknown message type supplied in message header. **System Action:** Request is failed but connection remains open. **Response:** Send a supported message type.
EINVAL(121)	NMsRsnExcessiveSize (4)	Excessive message size. **System Action:** Connection is closed. **Response:** Re-issue the connection and send a correctly formatted message.

Table 24. NSS XMLAppliance client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL(121)	NMsRsnHdrSize (5)	Message header size is not valid. **System Action:** Request is failed but connection remains open. **Response:** Send a message with the header size field set to the correct value.
EINVAL(121)	NMsRsnMsgSize (6)	Message size is not valid. For example, the message may be too short, or the message size may be greater than the sum of its parts. **System Action:** Request is failed but connection remains open. **Response:** Send a correctly formatted message.
EINVAL(121)	NMsRsnReservedNonzero (7)	Reserved data in message header, record header, or record data is non-zero value. Reserved fields must be set to 0 for compatibility with any future versions of the interface. **System Action:** Request is failed but connection remains open. **Response:** Send a message with reserved fields set to 0.
EINVAL(121)	NMsRsnRecordLength (8)	Unrecognized record length supplied in message. Length does not correspond to known record data. **System Action:** Request is failed but connection remains open. **Response:** Send a message with input filters of the correct length.
EINVAL(121)	NMsRsnRecordCount (9)	Unsupported record count supplied in message. NMI requests currently support a maximum of twenty input records. **System Action:** Request is failed but connection remains open. **Response:** Send a message with the correct number of input filters.
EINVAL(121)	NMsRsnSectionLength (10)	Unrecognized section length supplied in record. Length does not correspond to known section data. **System Action:** Request is failed but connection remains open. **Response:** Send a message with correct input filters.
EINVAL(121)	NMsRsnSectionCount (11)	Unrecognized section count supplied in record. NMI requests currently allow one section in an input record. **System Action:** Request is failed but connection remains open. **Response:** Send a message with correct input filters.
EINVAL(121)	NMsRsnClientAlreadyConnected (10002)	The remote client name is already registered with the NSS server. NSS client names must be unique. **System Action:** Connection is closed. **Response:** Re-issue the connection request using a unique client name.

Table 24. NSS XMLAppliance client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL(121)	NMsRsnNoMatchingCert (10004)	The NSS server could not find a matching certificate. The certificate does not exist on the NSS servers configured keyring, it is marked untrusted, or the NSS client does not have the authority to use the certificate. **System Action:** Request is failed but connection remains open. **Response:** Re-issue the request with an existing, trusted, and authorized certificate label.
EINVAL(121)	NMsRsnNoCertRep (10014)	The NSS server does not have a certificate repository available to process the request. **System Action:** Request is failed but connection remains open. **Response:** Start the appropriate application logic.
EINVAL(121)	NMsRsnInvalidService (10021)	A service has been requested that is not affiliated with the requested discipline. **System Action:** Connection is closed. **Response:** Re-attempt the connection and request only the services affiliated with the requested discipline.
EINVAL(121)	NMsRsnInvalidIdentity (10022)	The SAF user access check request did not contain a valid user identity. The SAF ID is not recognized or, if a certificate was provided as input to the NSS_CheckUserAccessReqToSrv, no valid certificate name filter mapping is defined in RACF. **System Action:** NSS server processing continues. **Response:** SEnsure that the user identity is entered correctly and reissue the request.
EINVAL(121)	NMsRsnInvalidSAFClass (10023)	The SAF class specified in the NSS_SAFCheckUserAccessReqToSrv message was unsupported. **System Action:** None. **Response:** Contact the NSS client vendor. The SERVAUTH class is the only currently-supported SAF class.
EINVAL(121)	NMsRsnInvalidProfileLength (10024)	The SAF user access check request contained an invalid profile length. **System Action:** NSS server processing continues. **Response:** Contact the NSS client vendor. The maximum profile length for the SERVAUTH class is 64 bytes.
EINVAL(121)	NMsRsnInvalidDiscipline (10025)	The discipline specified in the connection request contains an invalid value. **System Action:** Connection is closed. **Response:** Re-attempt the connection and pass in a valid discipline.

Table 24. NSS XMLAppliance client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL(121)	NMsRsnBadUpdate (10026)	The client has attempted to update its client information using values that cannot be changed after the initial connection has succeeded. **System Action:** Request is failed but connection remains open. **Response:** Re-attempt the update by changing only those fields which are acceptable under and update.
EINVAL(121)	NMsRsnInvalidAPIVersion (10027)	An NSS client has attempted to connect to the NSS server and has specified adherence to an API version that is insufficient for the requested discipline. **System Action:** Connection is closed. **Response:** Contact the NSS client vendor. NSS XMLAppliance clients must adhere to NMsec_NSS_API_VERSION2 (2) or higher.

EINVAL(121)	NMsRsnInvalidAccessLevel (10028)	The SAF user access check request contained an invalid value for the requested access level.
		System Action: NSS server processing continues.
		Response: Contact the NSS client vendor. Supported access levels are:
		`Requested access  Hexadecimal value`
		`READ  0x02`
		`UPDATE  0x04`
		`CONTROL  0x08`
		`ALTER  0x80`

EINVAL(121)	NMsRsnInvalidClientName (10029)	NSS_ConnectClientReqToSrv or NSS_UpdateClientInfoReqToSrv request is invalid. **System Action:** If the client name is invalid on the connect, the request is failed and the connection is closed. If the client name is invalid on the update, the request is failed, the connection remains open, but the client remains in the update pending state until a valid update is provided. **Response:** Re-attempt the connect or update by providing a valid NSS client name. Valid characters are [a-zA-Z0-9_-]. The client name must be left-justified and blank-padded. Embedded spaces are invalid.
EINVAL(121)	NMsRsnInvalidCertLabelName (10030)	The NSS XMLAppliance client request contained an invalid value for the requested certificate label name. **System Action:** NSS server processing continues. **Response:** The NSS client should re-issue the request with a valid certificate label name. The values accepted in this field are documented in the request input for the NSS_GetCertificateReqToSrv call.

Table 24. NSS XMLAppliance client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL(121)	NMsRsnNoPrivateKey (10031)	The certificate does not contain the private key. **System Action:** Request is failed but connection remains open. **Response:** If the certificate is intended to have a private key, then contact the NSSD administrator to determine what action to take.
EACCES(111)	NMsRsnDisconnectPending (1)	A client disconnect operation is pending, so no new request messages are being accepted. **System Action:** Request is rejected and connection is eventually closed. **Response:** Stop sending requests for this connection. Reconnect to the server and re-issue the request.
EACCES(111)	NMsRsnUpdatePending (2)	A client update operation is pending, so no new request messages are being accepted. **System Action:** None. **Response:** Re-issue the request.
EACCES(111)	NMsRsnNoAuthForService (4)	Userid is not authorized to use the NSS server for the requested service. **System Action:** Connection is closed. **Response:** Ensure that the clients access to the requested service is defined in the servers SERVAUTH profiles (EZB.NSS.sysname.clientname.discipline.service).
EACCES(111)	NMsRsnUserAuthentication (10001)	User authentication failed. **System Action:** NSS server processing continues. **Response:** Ensure that the userid is defined in the server's security manager and that the password or passticket is formed correctly.
EACCES(111)	NMsRsnSAFUserNotAuthenticated (10002)	The SAF user identity specified in the request failed authentication. **System Action:** The NSS server successfully completed the authentication check. Processing continues. **Response:** NSS client independent. The NSS client should react accordingly.
EACCES(111)	NMsRsnSAFUserAccessDenied (10003)	The SAF user access check indicates access denied from the security server. **System Action:** The NSS server successfully completed the access check. Processing continues. **Response:** NSS client independent. The NSS client should react accordingly.

Table 24. NSS XMLAppliance client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EACCES(111)	NMsRsnSAFResourceError (10004)	The SAF user access check indicates that a SAF server is not installed, has not been started, or the specified class is not active, is not defined, or that no profile exists for the specified resource. **System Action:** The NSS server failed to complete the access check. Processing continues. **Response:** NSS client independent. The NSS client should react accordingly. Most commonly this reason code indicates that the profile that was queried does not exist. It would then be up to the NSS client to decide whether this implies access denied or access granted.
EACCES(111)	NMsRsnUnsupportedDiscipline (10005)	The discipline specified in the connection request is currently disabled in the NSS server. **System Action:** Connection is closed. **Response:** Modify the NSS server configuration to enable the specified discipline.
EACCES(111)	NMsRsnNoAuthForPrivKey (10006)	Userid is not permitted to use the requested certificate's private key. **System Action:** Request is failed but connection remains open. **Response:** Permit user to security resource SERVAUTH profile EZB.NSSCERT. sysname.mappedlabelname. PRIVKEY
EACCES(111)	NMsRsnPrivKeyProtected (10007)	The certificate's private key is protected from being retrieved. **System Action:** Request is failed but connection remains open. **Response:** If the private key is intended to be retrieved, then contact the NSSD administrator to determine what action to take.
EACCES(111)	NMsRsnPrivKeyNotProtected (10008)	The certificate's private key is not stored in the ICSF PKA key data set (PKDS). Signature generation and decryption require use of certificates for which the private keys are stored in the ICSF PKDS. **System Action:** Request is failed but connection remains open. **Response:** If the private key is intended to be used for signature generation or decryption, then contact the NSSD administrator to determine what action to take.

Table 24. NSS XMLAppliance client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EACCES(111)	NMsRsnNoAuthForCert (10009)	Userid is not permitted to retrieve the requested certificate. **System Action:** Request is failed but connection remains open. **Response:** Send a message with correct input filters. Permit user to proper security resource SERVAUTH profiles: Read access to either of the following profiles will authorize the NSS server to retrieve the requested certificate: • EZB.NSSCERT. sysname.mappedlabelname. CERTAUTH • EZB.NSSCERT. sysname.mappedlabelname. HOST
ENOMEM(132)	0	Insufficient storage available in the server to process the request. **System Action:** Request is failed but connection remains open. **Response:** Increase the REGION size for the NSS server, or send a message with a narrower set of input filters to limit the response.
ENOLCK(131)	0	Failed to obtain an internal lock. **System Action:** Request fails but connection remains open. A message will appear in the MVS system log with additional diagnostic information. **Response:** Contact IBM service.
EGSKCMS(10004)	See gskcms.h. The Reason code represents the GSK (system SSL) return code provided on the failed call.	A System SSL error was encountered when issuing a System SSL library call. The NSS reason code will contain the System SSL CMS Status Code. **System Action:** Request is failed but connection remains open. **Response:** Review the system SSL CMS status codes from z/OS Cryptographic Services System SSL Programming.
ECSFBEXT(10005)	The high-order 16 bits of the reason code represent the ICSF return Code. The low-order 16 bits of the reason code represent the ICFS reason code.	An Integrated Cryptographic Service Facility (ICSF) error was encountered while performing an RSA operation. The reason code will contain the ICSF return code(high order 16 bits) and reason code(low order 16 bits). **System Action:** Request is failed but connection remains open. **Response:** Review the ICSF return and reason codes from z/OS Cryptographic Services ICSF Application Programmer's Guide.

Network security services server debug information

Obtaining syslog debug information for the network security service server

The SyslogLevel parameter in the network security services server configuration file controls the level of NSS server internal debug information that is sent to syslog. When configuring with the IBM Configuration Assistant for z/OS Communications Server, use the NSS Daemon Syslog Trace panel to configure the level of network security services server internal debug information that is sent to syslog. See z/OS Communications Server: IP Configuration Reference or the IBM Configuration Assistant for z/OS Communications Server online help for more information.

Abends

Messages and error-related information should be sent to the system console when an abend occurs during NSS server processing. NSSD will initiate a system dump for the abend condition. System dumps of the NSS server include Language Environment data. The Language Environment IPCS verbexit LEDATA can be used to format this information. See z/OS Language Environment Debugging Guide for more information. The following is a sample IPCS **verbexit ledata** command:

```
verbx ledata 'asid(68) tcb(007E5E88) ceedump nthreads(*)'
```

In this example, the network security services server asid is 0x68 and the address of the abended NSS TCB is 0x007E5E88.

Error codes

Several messages display a return code and reason generated by the NSS server. Most of these return codes and reasons are generated in support of the application interface for managing IP filtering and IPSec on remote network security clients. These return codes and reasons are documented in z/OS Communications Server: IP Programmer's Guide and Reference.

Additional return codes and reasons may be generated by the NSS server. These return codes and reasons are generated in support of remote management services offered to remote network security clients and are explained in the following error codes table.

Table 25. NSS IPSec client API return codes and reason codes

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EGSKSIGN (10001)	gsk_status code generated during the failure. Common codes are: **[CMSERR_ALG_NOT_SUPPORTED]** The signature algorithm is not supported. **[CMSERR_BAD_DIGEST_SIZE]** The certificate private key is not long enough to be used with the digest size required by the requested hash algorithm. **[CMSERR_KEY_MISMATCH]** The supplied key does not match the signature algorithm. **[CMSERR_NO_MEMORY]** Insufficient storage is available. **[CMSERR_ICSF_NOT_AVAILABLE]** ICSF is not available. **[CMSERR_ICSF_SERVICE_FAILURE]** An ICSF service failed.	A System SSL CMS error was encountered while attempting to create a signature. The reason code will contain the System SSL return code. **System Action** Request fails but connection remains open. **Response** Examine gsk_status code (returned as the reason code), which are documented in z/OS Cryptographic Services System SSL Programming. Verify the failed message contained correct data. If it did not then take action to correct the message content. If it did then contact the NSSD administrator to determine what action to take. If the gsk_status_code is CMS_ERR_ICSF_NOT_AVAILABLE, request that the NSSD administrator verify that ICSF is started. If the gsk_status_code is CMS_ERR_ICSF_SERVICE_FAILURE, notify the NSSD administrator. The NSSD administrator should determine whether the SAF CSFSERV general resource class is defined and determine whether the CSF1PKS profile is defined for that resource. If the CSF1PKS profile is defined, verify that NSSD has read access to it. See the z/OS Cryptographic Services ICSF Administrator's Guide for more information about the CSFSERV general resource and the CSF1PKS profile.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description

EGSKVAL (10002)	gsk_status code generated during the failure. Common codes are:
	[CMSERR_BAD_HANDLE]
	The database handle is not valid.
	[CMSERR_BAD_ISSUER_NAME]
	The certificate issuer name is not valid.
	[CMSERR_BAD_SIGNATURE]
	The signature is not correct.
	[CMSERR_CERT_CHAIN_NOT_TRUST]
	The certification chain is not trusted.
	[CMSERR_CERTIFICATE_REVOKED]
	The certificate is revoked.
	[CMSERR_EXPIRED]
	The certificate is expired.
	[CMSERR_INCORRECT_DBTYPE]
	The database type does not support certificates.
	[CMSERR_INCORRECT_KEY_USAGE]
	The issuer certificate does not allow signing certificates
	[CMSERR_ISSUER_NOT_CA]
	The certificate issuer is not a certification authority.
	[CMSERR_ISSUER_NOT_FOUND]
	The issuer certificate is not found in one of the data sources.
	[CMSERR_NAME_CONSTRAINTS_VIOLATED]
	The certificate name is not consistent with the name constraints.
	[CMSERR_NAME_NOT_SUPPORTED]
	The AuthorityKeyIdentifier extension name is not a directory name.
	[CMSERR_NOT_YET_VALID]
	The certificate is not yet valid.
	[CMSERR_PATH_TOO_LONG]
	The certification chain exceeds the maximum allowed by the CA.
	[CMSERR_SELF_SIGNED_NOT_FOUND]
	A self-signed certificate is not found in a trusted data source
	[CMSERR_ICSF_NOT_AVAILABLE]
	ICSF is not available.
		GSK validate certificate failure.
		System Action
		Request fails but connection remains open.
		Response
		Examine gsk_status code (returned as the reason code), which are documented in z/OS Cryptographic Services System SSL Programming. Verify the failed message contained correct data. If it did not, then take action to correct the message content. If it did and the reason code is one of the following, contact the certificate owner and inform them of the problem encountered with the certificate:
		`CMSERR_BAD_ISSUER_NAME`
		`CMSERR_BAD_SIGNATURE`
		`CMSERR_CERTIFICATE_REVOKED`
		`CMSERR_EXPIRED`
		`CMSERR_INCORRECT_KEY_USAGE`
		`CMSERR_ISSUER_NOT_CA`
		`CMSERR_NAME_CONSTRAINTS`
		`_VIOLATEDCMSERR_NAME_NOT_SUPPORTED`
		`CMSERR_NOT_YET_VALID`
		`CMSERR_PATH_TOO_LONG`
		If the reason code is anything other than the codes above, contact the NSSD administrator to determine what action to take. Other common reason codes include:
		`CMSERR_BAD_HANDLE`
		`CMSERR_CERT_CHAIN_NOT_TRUST`
		`CMSERR_INCORRECT_DBTYPE`
		`CMSERR_ISSUER_NOT_FOUND`
		`CMSERR_SELF_SIGNED_NOT_FOUND`
		If the gsk_status_code is CMSERR_ICSF_NOT_AVAILABLE, request that the NSSD administrator verify that ICSF is started.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EGSKVER (10003)	gsk_status code generated during the failure. Common codes are: **[CMSERR_ALG_NOT_SUPPORTED]** The signature algorithm is not supported. **[CMSERR_BAD_DIGEST_SIZE]** The digest size is not correct. **[CMSERR_BAD_SIGNATURE]** The signature is not correct. **[CMSERR_KEY_MISMATCH]** The supplied key does not match the signature algorithm. **[CMSERR_ICSF_NOT_AVAILABLE]** ICSF is not available. **[CMSERR_ICSF_SERVICE_FAILURE]** An ICSF service failed.	A System SSL CMS error was encountered while attempting to verify a signature. The reason code will contain the System SSL return code. **System Action** Request fails but connection remains open. **Response** Examine gsk_status code (returned as the reason code), which are documented in z/OS Cryptographic Services System SSL Programming. Verify the failed message contained correct data. If it did not, then take action to correct the message content. If it did, then treat the signature as an invalid signature. If the gsk_status_code is CMS_ERR_ICSF_NOT_AVAILABLE, request that the NSSD administrator verify that ICSF is started. If the gsk_status_code is CMS_ERR_ICSF_SERVICE_FAILURE, notify the NSSD administrator. The NSSD administrator should determine whether the SAF CSFSERV general resource class is defined and determine whether the CSF1PKV profile is defined for that resource. If the CSF1PKV profile is defined, verify that NSSD has read access to it. See z/OS Cryptographic Services ICSF Administrator's Guide for more information about the CSFSERV general resource and the CSF1PKV profile.
EGSKCMS (10004)	gsk_status code generated during the failure. Common codes are: **[CMSERR_ICSF_NOT_AVAILABLE]** ICSF is not available.	A System SSL CMS error was encountered while processing the request. The reason code will contain the System SSL return code. **System Action** Request fails but connection remains open. **Response** Examine gsk_status code (returned as the reason code), which are documented in z/OS Cryptographic Services System SSL Programming. If the gsk_status_code is CMS_ERR_ICSF_NOT_AVAILABLE, request that the NSSD administrator verify that ICSF is started.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
ECSFBEXT (10005)	The high-order 16 bits of the reason code represent the ICSF return code. The low-order 16 bits of the reason code represent the ICFS reason code.	An Integrated Cryptographic Service Facility (ICSF) error was encountered. The reason code will contain the ICSF return code (high-order 16 bits) and reason code (low-order 16 bits). **System Action** Request is failed but connection remains open. **Response** Review the ICSF return and reason codes from z/OS Cryptographic Services ICSF Application Programmer's Guide.
EACCES (111)	NMSRsnUserAuthentication (10001)	User authentication failed **System Action** Request fails and the connection is closed. **Response** Verify the following: The user ID under which the NSS client connects to the NSS server is correct The password used to authenticated that user ID is valid, or the application key used to generate the passticket is correct (this key is stored in the SAF-enabled security manager).
EACCES (111)	NMsRsnNoAuthForService (4)	The NSS client does not have access to the requested service through the governing SERVAUTH profile. **System Action** Request fails but connection remains open. **Response** If appropriate, define a SERVAUTH profile that will allow the requested access.
EACCES (111)	NMsRsnNoAuthForClientname (3)	The user ID in the connection request is not authorized to act on behalf of the NSS clientName. **System Action** Request fails and the connection is closed. **Response** Ensure that all of the following are correct: The user ID (and password, if necessary) as configured at the client. The client name as configured at the client. Also ensure that the appropriate SERVAUTH profiles are defined at the server system for the client.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EACCES (111)	NMsRsnDisconnectPending (1)	A disconnect operation is pending. **System Action** Request fails but connection remains open for a very short time. **Response** The client must reconnect the server before any more NSS services can be requested.
ECCESS (111)	NMsRsnUnsupportedDiscipline (10005)	The discipline specified in the connection request is currently disabled in the NSS server. **System Action** Connection is closed. **Response** Modify the NSS server configuration to enable the specified discipline.
EINVAL (121)	NMSRsnClientAlreadyConnected (10002)	Client is already connected to this server. **System Action** Request fails and the connection is closed. **Response** If appropriate, disconnect the active client and reattempt the connection request.
EINVAL (121)	NSSRsnRIDNotInCert (10003)	The certificate used to sign does not contain remote ID specified. **System Action** Request fails but connection remains open. **Response** None - this is an informational code only.
EINVAL (121)	NSSRsnBadCert (10005)	Certificate not valid. **System Action** Request fails but connection remains open. **Response** If the failing certificate is one that is stored on the local system, it should be refreshed or replaced. If that certificate comes from a remote system, then this is an informational code only.
EINVAL (121)	NSSRsnUnsupportedCert (10006)	Unsupported certificate encoding. **System Action** Request fails but connection remains open. **Response** Contact IBM service.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL (121)	NSSRsnBadLIDType (10007)	Unrecognized LID type. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnBadLIDValue (10008)	LID value not valid. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnBadRIDType (10009)	Unrecognized LID type. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnBadRIDValue (10010)	LID value not valid. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnBadLocalIPaddr (10011)	Local IPaddr not valid. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnBadRemoteIPaddr (10012)	Remote IPaddr not valid. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnAddrVersionMismatch (10013)	Local and remote IP address versions don't match. **System Action** Request fails but connection remains open. **Response** Contact IBM service.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL (121)	NSSRsnNoCertRep (10014)	Certificate repository not available. **System Action** Request fails but connection remains open. **Response** Create or restore the certificate repository and then try the request again.
EINVAL (121)	NSSRsnBadHashSize (10016)	Hash size not valid for specified hash algorithm. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnBadHashAlg (10017)	Hash algorithm not supported or an NSSD server is at a lower version than the IKED client. Before calling IBM service, check for msg EZD1904E in your IKED log. If it is a version mismatch, either change your IpSec policy to specify only algorithms that this version of NSSD supports or upgrade the NSSD server to the same version as IKED. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnSaNotInCertLife (10018)	SA lifetime not in certificate lifetime. **System Action** Request fails but connection remains open. **Response** None - this is an informational code only.
EINVAL (121)	NSSRsnBadCa (10019)	The DER encoding type specified for the Certificate Authority name is unrecognized. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnUnsupportedCaType (10020)	Unsupported CA encoding. **System Action** Request fails but connection remains open. **Response** Contact IBM service.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL (121)	NMsRsnInvalidService (10021)	A service has been requested that is not affiliated with the requested discipline. **System Action** Connection is closed. **Response** Re-attempt the connection and request only the services affiliated with the requested discipline.
EINVAL (121)	NMsRsnInvalidDiscipline (10025)	The discipline specified in the connection request contains an invalid value. **System Action** Connection is closed. **Response** Re-attempt the connection and pass in a valid discipline.
EINVAL (121)	NMsRsnBadUpdate (10026)	The client has attempted to update its client information using values that cannot be changed after the initial connection has succeeded. **System Action** Request is failed but connection remains open. **Response** Re-attempt the update by changing only those fields which are acceptable under an update.
EINVAL (121)	NMsRsnInvalidAPIVersion (10027)	An NSS client has attempted to connect to the NSS server and has specified adherence to an API version that is insufficient for the requested discipline. **System Action** Connection is closed. **Response** Re-attempt the connection using an accepted API version. NSS IPSec clients must adhere to NMsec_NSS_API_VERSION1 (1) or higher. NSS XMLAppliance clients must adhere to NMsec_NSS_API_VERSION2 (2) or higher.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL (121)	NMsRsnInvalidClientName (10029)	NSS_ConnectClientReqToSrv or NSS_UpdateClientInfoReqToSrv request is invalid. **System Action** If the client name is invalid on the connect, the request is failed and the connection is closed. If the client name is invalid on the update, the request is failed, the connection remains open, but the client remains in the update pending state until a valid update is provided. **Response** Re-attempt the connect or update by providing a valid NSS client name. Valid characters are [a-zA-Z0-9_-]. The client name must be left-justified and blank-padded. Embedded spaces are invalid.
EINVAL (121)	NSSRsnBadAuthMethod (10032)	Authentication method not supported. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnBadPRFAlg (10033)	PRF algorithm not supported. **System Action** Request fails but connection remains open. **Response** Contact IBM service.
EINVAL (121)	NSSRsnMissingCRLs (10034)	The NSS client requested strict revocation checking but certificate revocation lists (CRLs) are missing from the request. **System Action** Request fails but connection remains open. **Response** Re-attempt the request providing the missing CRLs.
EINVAL (121)	NSSRsnPRFAlgNotFIPS (10035)	The NSS server is configured for FIPS mode but the NSS client requested a PRF algorithm that is not valid for FIPS mode (e.g. HMAC-MD5). **System Action** Request fails but connection remains open. **Response** Re-attempt the request with a PRF algorithm that is valid for FIPS mode.

Table 25. NSS IPSec client API return codes and reason codes (continued)

Return code (NMsMRc)	Reason code (NMsMRsn)	Description
EINVAL (121)	NSSRsnHashAlgNotFIPS (10036)	The NSS server is configured for FIPS mode but the NSS client requested a hash algorithm that is not valid for FIPS mode (e.g. MD5). **System Action** Request fails but connection remains open. **Response** Re-attempt the request with a hash algorithm that is valid for FIPS mode.
ENOLCK (131)	0	Failed to obtain an internal lock. **System Action** Request fails but connection remains open. A message will appear in the MVS system log with additional diagnostic information. **Response** Contact IBM service.
ENOMEM (132)	NMsRsnTooManyConns (1)	The NSS server is already using its maximum number of 500 connections and cannot accept any more. **System Action** Connection is not opened and the request is failed. **Response** Try the request again later.
ENXIO (138)	NSSRsnUnknownClientName (10001)	The specified client name not recognized. **System Action** Request fails and the connection is closed. **Response** Verify that the client name was specified correctly and that the NSS client is connected to the NSS server. Note, however, that this error code often occurs when directing a request to an NSS client that is not currently connected to the NSS server.

TCP/IP services component trace for the network security services (NSS) server

The network security services (NSS) server uses component trace support to trace internal operations. This section describes how to specify NSS server trace and formatting options. For short descriptions of other tracing procedures, such as displaying trace status, see Chapter 5, "TCP/IP services traces and IPCS support," on page 45.

For detailed information, see the following information:

- z/OS MVS Diagnosis: Tools and Service Aids for information about component trace procedures.
- z/OS MVS Initialization and Tuning Reference for information about the component trace SYS1.PARMLIB member.
- z/OS MVS System Commands for information about commands.
- z/OS MVS Programming: Authorized Assembler Services Guide for procedures and return codes for component trace macros.
- z/OS MVS IPCS Commands for information about IPCS commands.
- z/OS MVS IPCS User's Guide for information about using IPCS.

Using CTRACE

You can specify component trace options at NSS server initialization or after the NSS server has initialized.

Table 26 lists the network security services server trace options.

Table 26. NSS server trace options

Trace events	Description
ALL	Select all types of records. **Note:** This option might have an impact on performance.
MINIMUM	Select the minimum level of tracing for the network security services server. This level includes the INIT, EXCEPT, and TERM categories.
INIT	Select NSS server initialization information.
TERM	Select NSS server termination information.
EXCEPT	Select NSS server exception information.
CONFIG	Select NSS server configuration information.
COMMANDS	Select processing of NSS server commands from the console or command line.
LOGMSGS	Select NSS server syslog messages. These entries can be used to easily correlate system log messages to a specific point in the CTRACE log.
ROUTING	Select NSS server threading and request dispatching information
SERIAL	Select NSS server serialization information
EVENT	Select NSS server event information.
SOCKETS	Select NSS server socket information.
PERFORM	Select NSS server performance information.
REQUESTS	Select NSS server request/response information.
FLOW	Select NSS server code flow information.
STORAGE	Select NSS server storage information.
CERTOPS	Select NSS server certificate operations information (cert cache ops and signature verify/create calls).
CONTROL	Select NSS server control information.

Table 26. NSS server trace options (continued)

Trace events	Description
DEBUG	Select NSS server debugging information.
VERBOSE	Select NSS server verbose debugging information.

Steps for enabling the CTRACE at network security service (NSS) server startup

A default minimum component trace is always started during NSS server initialization. Use a parmlib member to customize the parameters that are used to initialize the trace. The default NSS server component trace parmlib member is the SYS1.PARMLIB member CTINSS00. You can change the parmlib member name using the NSSD_CTRACE_MEMBER environment variable.

Rule: The NSS server reads the NSSD_CTRACE_MEMBER environment variable only during initialization. Changes to NSSD_CTRACE_MEMBER after server initialization have no effect.

For a description of trace options, see Table 26 on page 378.

Restriction: In addition to specifying the trace options, you can also change the NSS trace buffer size. The buffer size can be changed only at NSS initialization and has a maximum of 256 MB.

If the CTINSS00 member or the member that is specified in NSSD_CTRACE_MEMBER is not found when starting the network security services server, the following message is issued:

```
IEE5381 memberName MEMBER NOT FOUND IN PARMLIB
```

When this message is issued, the NSS component trace is started with a buffer size of 1 MB and the minimum tracing option.

Steps for enabling the CTRACE at network security services server startup

Edit the CTINSS00 parmlib member and enable the CTRACE at network security services server startup.

Procedure

Perform the following steps:

1. Edit the CTINSS00 parmlib member and specify the following:
   - TRACEOPTS ON
   - The buffer size you want with the BUFSIZE() parameter
   - The CTRACE options you want. To direct the CTRACE to an external writer, also specify the name of the writer JCL procedure in the WTR() parameter.

   See the example CTINSS00 parmlib member.
2. Start the network security services (NSS) server.

Steps for disabling the CTRACE at network security services server startup

Disable the CTRACE at network security services server startup.

Procedure

Follow these steps:

1. To disable the CTRACE at NSS startup, edit the CTINSS00 parmlib member and specify TRACEOPTS OFF.
2. Start the network security services (NSS) server.

Step for enabling the CTRACE after the network security services server is started

After the network security services (NSS) server is started, enable the CTRACE to an internal buffer or to an external writer.

Procedure

Perform the appropriate step:

- Issue the following console commands to enable the CTRACE to an internal buffer:

```
TRACE CT,ON,COMP=SYSTCPNS,SUB=(nss_jobname)
R xx,OPTIONS=(option[,option2...]),END
```

- Issue the following console commands to enable the CTRACE to an external writer:

```
TRACE CT,WTRSTART=writer_proc
TRACE CT,ON,COMP=SYSTCPNS,SUB=(nss_jobname)
R xx,OPTIONS=(option[,option2...]),WTR=writer_proc,END
```

Step for disabling the CTRACE after the network security services server is started

After the network security services (NSS) server is started, disable the CTRACE to an internal buffer or to an external writer.

Procedure

Perform the appropriate step:

- Issue the following console command to disable the CTRACE to an internal buffer:

```
TRACE CT,OFF,COMP=SYSTCPNS,SUB=(nss_jobname)
```

- Issue the following console commands to disable a CTRACE to an external writer:

```
TRACE CT,OFF,COMP=SYSTCPNS,SUB=(nss_jobname)
TRACE CT,WTRSTOP=writer_proc
```

Step for displaying the CTRACE status

To display the CTRACE status, issue the following console command:

```
D TRACE,COMP=SYSTCPNS,SUB=(nss_jobname)
```

Enabling CTRACE after network security services server initialization

After NSS initialization, you must use the TRACE CT command to change the component trace options. Each time a new component trace is initiated, all prior trace options are turned off and the new options are put into effect. You can specify the trace options with or without the parmlib member. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45.

Formatting network security services server trace records

You can format component trace records using IPCS panels or a combination of the IPCS panels and the CTRACE command, either from a dump or from external-writer files. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45.

Enter any combination of values as options to filter the CTRACE entries. The options must be entered using the following format:

```
TYPE(option[,option]...)
```

You can use any of the options listed in Table 26 on page 378, except ALL and MINIMUM.

Chapter 11. Diagnosing dynamic VIPA and sysplex problems

This topic presents diagnostic information for dynamic virtual IP address (DVIPA) and sysplex problems, and contains the following subtopics:

- "Overview of diagnosing sysplex distributor problems"
- "Steps for diagnosing sysplex problems" on page 384
- "Steps for diagnosing problems using DVIPAs in source IP address selection for TCPconnections problems" on page 396
- "Steps for diagnosing problems with the SYSPLEX-wide ephemeral port assignment for distributed DVIPAs" on page 397
- "Diagnosing problems with the SYSPLEX-wide ephemeral port assignment for EXPLICITBINDPORTRANGE processing" on page 399
- "Diagnosing sysplex-wide security association (SWSA) problems" on page 403
- "Steps for diagnosing sysplex routing problems" on page 409
- "Steps for diagnosing Tier 1 z/OS sysplex distribution problems" on page 411
- "Steps for diagnosing Tier 1 non-z/OS sysplex distribution problems" on page 413

Overview of diagnosing sysplex distributor problems

Diagnosing sysplex distributor problems presents some unique challenges. Because a DVIPA can be associated with multiple stacks in a sysplex, determining where a problem is can be more difficult. You can use a combination of the Netstat command from the system console and display sysplex commands to provide a clear picture of the sysplex. See z/OS Communications Server: IP Configuration Guide for an introduction to sysplex distribution with virtual addressing.

You can collect Netstat information in the following ways:

- You can issue the z/OS UNIX **netstat** command from the z/OS UNIX shell.
- You can issue the NETSTAT command from TSO.
- You can issue the DISPLAY TCPIP,,NETSTAT command from the system console.

In the following list of activities, you can find steps to perform them in "Steps for diagnosing sysplex problems" on page 384:

- First, determine that all the stacks that you expect to be communicating are in the same subplex, if subplexing is being used. See step 1 on page 384
- For problems where the actual DVIPAs defined on a stack are not what you expected, confirm the current definitions on a stack. See step 2 on page 385.
- For Sysplex Distributor workload monitoring, use steps 7 on page 390 and 10 on page 395. If the output from these commands is not what you expected, see step 6 on page 389 for an overall picture of all DVIPA activity in your sysplex.
- If the output from step 6 on page 389 reveals an expected target stack not listed for a distributed DVIPA, perform step 3 on page 386 on the target stack in question. This helps to identify configuration problems on that stack. Note what is required of target stacks. Also use step 11 on page 395 to verify that a server application has indeed been activated and bound to the correct port.

© Copyright IBM Corp. 2000, 2015

- To help follow the flow of packets into and throughout the sysplex, a CTRACE with options XCF, TCP, and SYSTCPDA on participating stacks is useful. Use these to:
  - Identify the connection being received by the distributing stack
  - Determine the stack to which the connection is forwarded
  - Verify the connection being forwarded
  - Determine the expected target stack receiving and processing the connection

  After the connection has been established, subsequent packets can be followed in the same manner. When the connection is terminated, CTRACE records record target stacks, cleans up the connection, and notifies the distributing stack.

Steps for diagnosing sysplex problems

Use these steps to diagnose sysplex problems. The output is shown in the long, or IPv6-enabled, format.

Procedure

Perform the following steps:

1. If subplexing is being used, run the D XCF,GROUP MVS command to determine what groups are being used in the sysplex. Find all the group names with the format EZBT*vvtt*.

```
D XCF,GROUP
IXC331I  10.11.09  DISPLAY XCF 637
   GROUPS(SIZE):   COFVLFNO(2)       EZBT11CS(3)       EZBT1121(2)
                   EZBT1122(2)       EZBT1123(1)       ISTCFS11(2)
                   ISTXCF11(2)       MVSVIC02(1)       MVSVIC11(1)
                   MVSVIC94(1)       MVSVIC96(1)       MVS031(1)
                   MVS165(1)         SYSDAE(9)         SYSENF(2)
                   SYSGRS(2)         SYSIEFTS(2)       SYSIGW00(9)
                   SYSIGW01(9)       SYSIKJBC(2)       SYSIOS01(2)
                   SYSJES(2)         SYSMCS(7)         SYSMCS2(4)
                   SYSTTRC(2)        SYSWLM(2)         VIC140(1)
```

   Run the D XCF,GROUP,groupname MVS command for each of the EZBT*vvtt* format groups to find all the member stacks participating in each subplex group.

```
D XCF,GROUP,EZBT11CS
IXC332I  10.11.18  DISPLAY XCF 640
 GROUP EZBT11CS:  MVS165TCPCS1      MVS165TCPCS3      VIC011TCPCS2
D XCF,GROUP,EZBT1121
IXC332I  10.11.18  DISPLAY XCF 640
 GROUP EZBT1121:  MVS165TCPCS11     VIC011TCPCS21
D XCF,GROUP,EZBT1122
IXC332I  10.11.18  DISPLAY XCF 640
 GROUP EZBT1122:  MVS165TCPCS12     VIC011TCPCS22
D XCF,GROUP,EZBT1123
IXC332I  10.11.18  DISPLAY XCF 640
 GROUP EZBT1123:  VIC011TCPCS23
```

   The member names listed are the MVS name concatenated with the TCP/IP stack name. Verify that all the TCP/IP stacks you expect to be communicating are all within the same subplex group.

   If the TCP/IP stacks that you expect to be communicating use the same HiperSockets CHPID, verify that they have all specified the same IQDVLANID value. Issue the Netstat CONFIG/-f command for each stack and verify that the IQDVLANID value displayed in the Global Configuration Information section of the output is the same for each stack.

```
NETSTAT CONFIG
 MVS TCP/IP NETSTAT CS V1R8       TCPIP Name: TCPCS           12:55:20
 TCP Configuration Table:
 DefaultRcvBufSize:  00016384  DefaultSndBufSize: 00016384
 DefltMaxRcvBufSize: 00262144
 MaxReTransmitTime:  120.000   MinReTransmitTime: 0.500
 RoundTripGain:      0.125     VarianceGain:      0.250
 VarianceMultiplier: 2.000     MaxSegLifeTime:    30.000
 DefaultKeepALive:   00000120  DelayAck:          Yes
 RestrictLowPort:    Yes       SendGarbage:       No
 TcpTimeStamp:       Yes       FinWait2Time:      600
 TTLS                Yes
 UDP Configuration Table:
 DefaultRcvBufSize: 00065535  DefaultSndBufSize: 00065535
 CheckSum:          Yes
 RestrictLowPort:   Yes       UdpQueueLimit:     No

 IP Configuration Table:
 Forwarding: Yes    TimeToLive: 00064  RsmTimeOut:  00060
 IpSecurity: Yes
 ArpTimeout: 01200  MaxRsmSize: 65535  Format:      Short
 IgRedirect: Yes    SysplxRout: No     DoubleNop:   No
 StopClawEr: No     SourceVipa: No
 MultiPath:  Conn   PathMtuDsc: No     DevRtryDur:  0000000090
 DynamicXCF: Yes
   IpAddr/PrefixLen: 193.9.200.3/28    Metric: 01
   SecClass: 8
 IQDIORoute: Yes       QDIOPriority: 1
 TcpStackSrcVipa: 201.1.10.10

 SMF Parameters:
 Type 118:
   TcpInit:      01   TcpTerm:    02   FTPClient:    03
   TN3270Client: 00   TcpIpStats: 05
 Type 119:
   TcpInit:      No   TcpTerm:    No   FTPClient:    Yes
   TcpIpStats:   Yes  IfStats:    Yes  PortStats:    Yes
   Stack:        Yes  UdpTerm:    Yes  TN3270Client: Yes

 Global Configuration Information:
 TcpIpStats: Yes  ECSALimit: 0002047K  PoolLimit: 2096128K
 MlsChkTerm: No   XCFGRPID:  11   IQDVLANID: 27
 Sysplex Monitor:
   TimerSecs:  60    Recovery:  Yes    DelayJoin:  No   AutoRejoin: Yes
 MONINTF: NO DYNROUTE: NO
 Network Monitor Configuration Information:
 PktTrcSrv: Yes  TcpCnnSrv: Yes   MinLifTim: 3   SmfSrv: Yes

 Data Trace Setting:
 JobName: *           TrRecCnt: 00000009  Length: FULL
 IpAddr:  *                SubNet: *
```

2. Run the Netstat VIPADCFG/-F display command on the distributing stack to confirm that it is configured to distribute the DVIPA and how it is to be distributed. If the DVIPA has been deactivated, the deactivated configuration definitions are displayed under the heading DEACTIVATED DYNAMIC VIPA INFORMATION.
   - Figure 31 on page 386 shows that the TCP/IP identified by TCPCS was configured to distribute DVIPAs. Workload for the first DVIPA, 201.2.10.11 ports 20 and 21, is being distributed to all stacks in the sysplex including TCPCS itself; the configured distribution method is SERVERWLM.
   - Workload for 201.2.10.12, ports 20 and 21, is being distributed only to the TCP/IP with dynamic XCF address 193.9.200.2.

- Workload for 201.2.10.13 port 5000 is being distributed to all stacks using the TIMEDAFFinity function.
- Workload for IPv6 DVIPA 2001:0DB8:1::1, port 6000 is being distributed to all stacks; the configured distribution method is SERVERWLM.
- The DVIPA, 201.2.10.23, port 4000, was configured to be distributed to all stacks in the sysplex. Because the DVIPA has been deactivated on this stack, it is not currently being distributed by this stack.

```
D TCPIP,TCPCS,NET,VIPADCFG
EZD0101I NETSTAT CS V1R7 TCPCS 876
DYNAMIC VIPA INFORMATION:
  VIPA BACKUP:
    IPADDR/PREFIXLEN: 201.2.10.21
      RANK: 000080  MOVEABLE:            SRVMGR:
    IPADDR/PREFIXLEN: 201.2.10.22
      RANK: 000080  MOVEABLE:            SRVMGR:
  VIPA DEFINE:
    IPADDR/PREFIXLEN: 201.2.10.11/28
      MOVEABLE: IMMEDIATE  SRVMGR: NO
    IPADDR/PREFIXLEN: 201.2.10.12/28
      MOVEABLE: IMMEDIATE  SRVMGR: NO
    IPADDR/PREFIXLEN: 201.2.10.13/28
      MOVEABLE: IMMEDIATE  SRVMGR: NO
    INTFNAME: DVIPA1
      IPADDR: 2001:0DB8:1::1
        MOVEABLE: IMMEDIATE  SRVMGR: N/A
  VIPA DISTRIBUTE:
    DEST:        201.2.10.11..20
      DESTXCF:   ALL
        SYSPT:   NO   TIMAFF: NO    FLG: SERVERWLM
    DEST:        201.2.10.11..21
      DESTXCF:   ALL
        SYSPT:   NO   TIMAFF: NO    FLG: SERVERWLM
    DEST:        201.2.10.12..20
      DESTXCF:   193.9.200.2
        SYSPT:   NO   TIMAFF: NO    FLG: BASEWLM
    DEST:        201.2.10.12..21
      DESTXCF:   193.9.200.2
        SYSPT:   NO   TIMAFF: NO    FLG: BASEWLM
    DEST:        201.2.10.13..5000
      DESTXCF:   ALL
        SYSPT:   NO   TIMAFF: 45    FLG: BASEWLM
    DESTINTF:    DVIPA1
      DEST:      2001:0DB8:1::1..6000
        DESTXCF: ALL
          SYSPT: NO   TIMAFF: NO    FLG: SERVERWLM
Deactivated Dynamic VIPA Information:
    VIPA Define:
       IpAddr/PrefixLen: 201.2.10.23/28
         Moveable: Immediate  SrvMgr: No
    VIPA Distribute:
       Dest:        201.2.10.23..4000
         DestXCF:   ALL
           SysPt:   No   TimAff: No    Flg: BaseWLM
```

Figure 31. Netstat VIPADCFG/-F example

3. Run the display Netstat CONFIG/-f command on the distributing stack and all target stacks to confirm that the correct IPCONFIG and IPCONFIG6 options have been specified.

   Specify SYSPLEXROUTING on the distributor and all target stacks in order to get WLM-based distribution. Verify that DYNAMICXCF was specified on the distributor and all target stacks.

Figure 32 shows the output of this command for the distributing TCP/IP:

```
D TCPIP,TCPCS,N,CONFIG
EZD0101I NETSTAT CS V1R8 TCPCS 928
TCP CONFIGURATION TABLE:
DEFAULTRCVBUFSIZE:  00016384  DEFAULTSNDBUFSIZE: 00016384
DEFLTMAXRCVBUFSIZE: 00262144
MAXRETRANSMITTIME:  120.000   MINRETRANSMITTIME: 0.500
ROUNDTRIPGAIN:      0.125     VARIANCEGAIN:      0.250
VARIANCEMULTIPLIER: 2.000     MAXSEGLIFETIME:    30.000
DEFAULTKEEPALIVE:   00000120  DELAYACK:          YES
RESTRICTLOWPORT:    YES       SENDGARBAGE:       NO
TCPTIMESTAMP:       YES       FINWAIT2TIME:      600
UDP CONFIGURATION TABLE:
DEFAULTRCVBUFSIZE: 00065535  DEFAULTSNDBUFSIZE: 00065535
CHECKSUM:          YES
RESTRICTLOWPORT:   YES       UDPQUEUELIMIT:     YES
IP CONFIGURATION TABLE:
FORWARDING: YES    TIMETOLIVE: 00064  RSMTIMEOUT:  00060
ARPTIMEOUT: 01200  MAXRSMSIZE: 65535  FORMAT:      LONG
IGREDIRECT: YES    SYSPLXROUT: YES    DOUBLENOP:   NO
STOPCLAWER: NO     SOURCEVIPA: YES
MULTIPATH:  NO     PATHMTUDSC: NO     DEVRTRYDUR:  0000000090
DYNAMICXCF: YES
  IPADDR/PREFIXLEN: 193.15.1.1/24       METRIC: 02
IQDIOROUTE: NO
TCPSTACKSRCVIPA: 203.15.1.1
IPV6 CONFIGURATION TABLE:
FORWARDING:    YES  HOPLIMIT:   00255  IGREDIRECT:  YES
SOURCEVIPA:    YES  MULTIPATH:  NO     ICMPERRLIM:  00003
IGRTRHOPLIMIT: NO
DYNAMICXCF: YES
  IPADDR: 2001:0DB8::151:0
  INTFID: 0006:0007:0008:0009
TCPSTACKSRCVIPA: DVIPA1
SMF PARAMETERS:
TYPE 118:
  TCPINIT:      00   TCPTERM:    00   FTPCLIENT:    00
  TN3270CLIENT: 00   TCPIPSTATS: 00
TYPE 119:
  TCPINIT:      NO   TCPTERM:    NO   FTPCLIENT:    NO
  TCPIPSTATS:   NO   IFSTATS:    NO   PORTSTATS:    NO
  STACK:        NO   UDPTERM:    NO   TN3270CLIENT: NO
GLOBAL CONFIGURATION INFORMATION:
TCPIPSTATS: NO   ECSALIMIT: 0000000K  POOLLIMIT: 0000000K
MLSCHKTERM: NO
SYSPLEX MONITOR:
  TIMERSECS: 0060  RECOVERY: YES  DELAYJOIN: NO   AUTOREJOIN: YES
  MONINTF:   NO    DYNROUTE: NO
```

Figure 32. Netstat CONFIG/-f example

Run the D WLM,SYSTEMS command on the distributing stack and all targets stack to confirm that WLM is active. For more information about the DISPLAY command, see z/OS MVS System Commands. Figure 33 on page 388 shows an example:

```
D WLM,SYSTEMS
IWM025I  16.38.58  WLM DISPLAY 963
  ACTIVE WORKLOAD MANAGEMENT SERVICE POLICY NAME: DEFAULT
  ACTIVATED: 2003/10/29  AT: 14:51:50  BY: N/A       FROM: VIC015
  DESCRIPTION: IBM'S WLM DEFAULT POLICY
  RELATED SERVICE DEFINITION NAME: N/A
  INSTALLED: 2003/10/29  AT: 14:51:50  BY: N/A       FROM: N/A
  WLM VERSION LEVEL:       LEVEL013
  WLM FUNCTIONALITY LEVEL: LEVEL001
  STRUCTURE SYSZWLM_WORKUNIT STATUS: DISCONNECTED
  *SYSNAME*  *MODE*  *POLICY*  *WORKLOAD MANAGEMENT STATUS*
  VIC015     GOAL    DEFAULT   ACTIVE, NOT RUNNING WITH ACTIVE POLICY
```

Figure 33. D WLM,SYSTEMS example

4. Run the display Netstat VIPADYN/-v command on the distributing stack to verify that the DVIPA status is ACTIVE and the distribution status is DIST or DIST/DEST. The deactivated DVIPA 203.2.10.23 do not appear in this display. Figure 34 shows an example:

```
D TCPIP,TCPCS,NET,VIPADYN
EZD0101I NETSTAT CS V1R8 TCPCS 827
IPADDR/PREFIXLEN: 201.2.10.11/28
  STATUS:  ACTIVE     ORIGIN: VIPADEFINE       DISTSTAT: DIST/DEST
ActTime: 03/02/2005 16:45:20
IPADDR/PREFIXLEN: 201.2.10.12/28
  STATUS:  ACTIVE     ORIGIN: VIPADEFINE       DISTSTAT: DIST
ActTime: 03/02/2005 16:45:20
IPADDR/PREFIXLEN: 201.2.10.13/28
  STATUS:  ACTIVE     ORIGIN: VIPADEFINE       DISTSTAT: DIST/DEST
ActTime: 03/02/2005 16:45:20
IPADDR/PREFIXLEN: 201.2.10.21
  STATUS:  BACKUP     ORIGIN: VIPABACKUP       DISTSTAT:
ActTime: n/a
IPADDR/PREFIXLEN: 201.2.10.22
  STATUS:  BACKUP     ORIGIN: VIPABACKUP       DISTSTAT:
ActTime: n/a
INTFNAME: DVIPA1
  IPADDR: 2001:0DB8:1::1
  STATUS:  ACTIVE     ORIGIN: VIPADEFINE       DISTSTAT: DIST/DEST
ActTime: 03/02/2005 16:45:20
```

Figure 34. Netstat VIPADYN/-v example

5. Run display Netstat VIPADYN/-v command on the target stacks to verify that they have activated the distributed DVIPA and have it designated as a DEST. In this case, TCPCS2 has designated the distributed DVIPAs as DEST and TCPCS2 is a backup stack for several DVIPAs (status and origin show backup). Figure 35 on page 389 shows an example:

```
D TCPIP,TCPCS2,NET,VIPADYN
EZD0101I NETSTAT CS V1R8 TCPCS2 905
IPADDR/PREFIXLEN: 201.2.10.11/28
  STATUS: BACKUP     ORIGIN: VIPABACKUP       DISTSTAT: DEST
  ActTime: 03/02/2005 16:45:20
IPADDR/PREFIXLEN: 201.2.10.12/28
  STATUS: BACKUP     ORIGIN: VIPABACKUP       DISTSTAT: DEST
  ActTime: 03/02/2005 16:45:20
IPADDR/PREFIXLEN: 201.2.10.13/28
  STATUS: ACTIVE     ORIGIN:                  DISTSTAT: DEST
  ActTime: 03/02/2005 16:45:20
IPADDR/PREFIXLEN: 201.2.10.21
  STATUS: BACKUP     ORIGIN: VIPABACKUP       DISTSTAT:
  ActTime: n/a
IPADDR/PREFIXLEN: 201.2.10.22
  STATUS: BACKUP     ORIGIN: VIPABACKUP       DISTSTAT:
  ActTime: n/a
INTFNAME: DVIPA1
  IPADDR: 2001:0DB8:1::1
    STATUS: ACTIVE     ORIGIN:                  DISTSTAT: DEST
    ActTime: 03/02/2005 16:45:20
```

Figure 35. Netstat VIPADyn/-v example

6. Run the Sysplex VIPADyn command from any stack in the sysplex to get a global view of how and where DVIPAs are defined within the sysplex and what their status is on each stack. Deactivated DVIPA configurations do not appear in this display. Figure 36 on page 390 shows the following:
   - Which TCP/IPs own distributed DVIPAs, DIST field=BOTH or DIST
   - Which TCP/IPs have been made targets, DIST field = DEST
   - The status of all other DVIPAs in this sysplex

```
D TCPIP,TCPCS2,SYSPLEX,VIPAD
EZZ8260I SYSPLEX CS V1R7 948
VIPA DYNAMIC DISPLAY FROM TCPCS2   AT VIC015
IPADDR: 201.2.10.11
  ORIGIN: VIPABACKUP
  TCPNAME  MVSNAME  STATUS RANK DIST
  -------- -------- ------ ---- ----
  TCPCS    VIC015   ACTIVE      BOTH
  TCPCS2   VIC015   BACKUP 100  DEST
  TCPCS3   VIC015   BACKUP 010  DEST
IPADDR: 201.2.10.12
  ORIGIN: VIPABACKUP
  TCPNAME  MVSNAME  STATUS RANK DIST
  -------- -------- ------ ---- ----
  TCPCS    VIC015   ACTIVE      DIST
  TCPCS2   VIC015   BACKUP 075  DEST
  TCPCS3   VIC015   BACKUP 010
IPADDR: 201.2.10.13
  TCPNAME  MVSNAME  STATUS RANK DIST
  -------- -------- ------ ---- ----
  TCPCS    VIC015   ACTIVE      BOTH
  TCPCS3   VIC015   BACKUP 010  DEST
  TCPCS2   VIC015   ACTIVE      DEST
IPADDR: 201.2.10.21
  ORIGIN: VIPABACKUP
  TCPNAME  MVSNAME  STATUS RANK DIST
  -------- -------- ------ ---- ----
  TCPCS3   VIC015   ACTIVE
  TCPCS2   VIC015   BACKUP 100
  TCPCS    VIC015   BACKUP 080
IPADDR: 201.2.10.22
  ORIGIN: VIPABACKUP
  TCPNAME  MVSNAME  STATUS RANK DIST
  -------- -------- ------ ---- ----
  TCPCS3   VIC015   ACTIVE
  TCPCS    VIC015   BACKUP 080
  TCPCS2   VIC015   BACKUP 075
INTFNAME: DVIPA1
IPADDR: 2001:0DB8:1::1
  TCPNAME  MVSNAME  STATUS RANK DIST
  -------- -------- ------ ---- ----
  TCPCS    VIC015   ACTIVE      BOTH
  TCPCS2   VIC015   ACTIVE      DEST
```

Figure 36. Sysplex VIPADyn example

7. Run the Netstat VDPT/-O command on the distributing stack to confirm that there are target stacks available with server applications ready.

   With the keyword DETAIL you can also see the following fields. For a complete DETAIL display example, see z/OS Communications Server: IP System Administrator's Commands.

   - Raw (before normalization) WLM composite weight
   - The original CP, zAAP, and zIIP weights, and the proportioned CP, zAAP, and zIIP weights that were used to determine the raw WLM composite weight. The original CP, zAAP, and zIIP weights, and the proportioned CP, zAAP, and zIIP weights are for only SERVERWLM and BASEWLM distribution algorithms.
   - The target server connection responsiveness factors that make up the TSR and the current WLM or QOS weight for each service level mapping to a DVIPA or port entry for each target stack (each DESTXCF ADDR). See the z/OS Communications Server: IP User's Guide and Commands for more information.

This display shows only target stacks that are currently up and have joined the sysplex. If there are fewer entries than what resulted from the display **d tcpip,,net,vipadcfg** command, the missing entries might be for target stacks that are not yet up, or for stacks that are already up now, but that do not specify the expected dynamic XCF address. Figure 37 on page 393 shows an example:

```
D TCPIP,TCPCS,NET,VDPT,DETAIL

 EZD0101I NETSTAT CS V1R7 TCPCS 010
 DYNAMIC VIPA DESTINATION PORT TABLE:
 DEST:201.2.10.11..20
 DESTXCF:193.9.200.2
 TOTALCONN:0000000959 RDY:001 WLM:11 TSR:   79
 FLG:SERVERWLM
     TCSR: 100 CER:  75 SEF: 79
     Weight 44
       Raw          CP: 44 zAAP: 00 zIIP: 00
       Proportional CP: 44 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01

 DEST:201.2.10.11..20
 DESTXCF:193.15.1.1
 TOTALCONN:0000000330 RDY:001 WLM:9 TSR:   68
 FLG:SERVERWLM
     TCSR:  87 CER:  75 SEF: 79
     Weight 36
       Raw          CP: 50 zAAP: 00 zIIP: 34
       Proportional CP: 05 zAAP: 00 zIIP: 31
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.11..20
 DESTXCF:193.15.3.1
 TOTALCONN:0000000315 RDY:001 WLM:15 TSR:  100
 FLG:SERVERWLM
     TCSR: 100 CER: 100 SEF: 100
     Weight 60
       Raw          CP: 60 zAAP: 00 zIIP: 00
       Proportional CP: 60 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.11..21
 DESTXCF:193.9.200.2
 TOTALCONN:0000000021 RDY:001 WLM:15 TSR:  100
 FLG:SERVERWLM
     TCSR: 100 CER: 100 SEF: 100
     Weight 60
       Raw          CP: 60 zAAP: 00 zIIP: 00
  Proportional CP: 60 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.11..21
 DESTXCF:193.15.1.1
 TOTALCONN:0000000008 RDY:001 WLM:11 TSR:   78
 FLG:SERVERWLM
     TCSR:  99 CER:  99 SEF:  80
     Weight 44
       Raw          CP: 44 zAAP: 00 zIIP: 00
       Proportional CP: 44 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.11..21
 DESTXCF:193.15.3.1
 TOTALCONN:0000000007 RDY:001 WLM:10 TSR:  94
```

```
FLG:SERVERWLM
     TCSR:  97 CER:  98 SEF: 97
     Weight 40
       Raw          CP: 40 zAAP: 00 zIIP: 00
       Proportional CP: 40 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.12..20
 DESTXCF:193.9.200.2
 TOTALCONN:0000000000 RDY:001 WLM:03 TSR:  99
 FLG:BASEWLM
     TCSR: 100 CER:  99 SEF:  99
     Weight 12
       Raw          CP: 20 zAAP: 11 zIIP: 00
  Proportional CP: 02 zAAP: 10 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.12..21
 DESTXCF:193.9.200.2
 TOTALCONN:0000000000 RDY:001 WLM:03 TSR: 100
 FLG:BASEWLM
     TCSR: 100 CER: 100 SEF: 100
     Weight 12
       Raw          CP: 12 zAAP: 00 zIIP: 00
       Proportional CP: 12 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.13..5000
 DESTXCF:193.9.200.2
 TOTALCONN:0000000000 RDY:001 WLM:03 TSR:   0
 FLG:BASEWLM
     TCSR:  90 CER:  75 SEF:   0
     Weight 12
       Raw          CP: 12 zAAP: 00 zIIP: 00
       Proportional CP: 12 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.13..5000
 DESTXCF:193.15.1.1
 TOTALCONN:0000000000 RDY:001 WLM:01 TSR:   27
 FLG:BASEWLM
     TCSR: 100 CER:  27 SEF: 27
     Weight 04
       Raw          CP: 04 zAAP: 00 zIIP: 00
       Proportional CP: 04 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DEST:201.2.10.13..5000
 DESTXCF:193.15.3.1
 TOTALCONN:0000000000 RDY:001 WLM:01 TSR:   48
 FLG:BASEWLM
     TCSR:  75 CER:  64 SEF:  64
     Weight 04
       Raw          CP: 04 zAAP: 00 zIIP: 00
       Proportional CP: 04 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
        W/Q:01
 DESTINTF:DVIPA1
 DEST:1::1..6000
 DESTXCF:FEC0::151:0
 TOTALCONN:0000000511 RDY:001 WLM:11 TSR:   79
  FLG:SERVERWLM
     TCSR:  99 CER:  98 SEF:  80
     Weight 44
       Raw          CP: 44 zAAP: 00 zIIP: 00
       Proportional CP: 44 zAAP: 00 zIIP: 00
      QosPlcAct:*DEFAULT*
```

```
       W/Q:01
DESTINTF:DVIPA1
DEST:1::1..6000
DESTXCF:FEC0::152:0
TOTALCONN:0000001410 RDY:001 WLM:10 TSR:    100
FLG:SERVERWLM
    TCSR: 100 CER:   95 SEF: 100
    Weight 40
      Raw          CP: 40 zAAP: 00 zIIP: 00
      Proportional CP: 40 zAAP: 00 zIIP: 00
     QosPlcAct:*DEFAULT*
       W/Q:01
```

Figure 37. Netstat VDPT/-O example

8. Examine the READY (RDY) count fields. The READY (RDY) count is the number of servers that are currently listening on the DVIPA and PORT specified in the DEST: field on the target stack that was identified by the DESTXCF address.

   For servers that use more than one port, the RDY value reflects the port where a LISTEN is performed. For example, for FTP, the control connection port (port 21) is where the RDY count is greater than 0. If the ready count is not as expected, proceed to step 11 on page 395 to verify whether any non-quiesced server is listening on the DPORT on the target stack. If there is a server listening on the target stack, verify that it has not been quiesced by a VARY TCPIP,,SYSPLEX,QUIESCE command. On the target stack, run the Netstat ALL/-A command and verify that the quiesced value is NO.
9. Check the TotalConn count to see the distribution history. This is a cumulative count of the number of connections that have been forwarded by the distributing stack to each target stack.

   If the connections are not being distributed to the target stacks as expected and either the WLM field or the W/Q field contains 00, then consider the following:

   - If using WLM to distribute connections based upon the workload of the target stacks, verify that all participating stacks (the distributor and all targets) have SYSPLEXROUTING specified. See step 3 on page 386 for instructions for verifying this. Also, verify that WLM is configured and active on all participating stacks. See step 3 on page 386.
   - If the WLM configuration appears correct and BASEWLM is being used as the distribution method, consider whether the unexpected distribution results might be caused by the current workload on the target stacks.

     If SERVERWLM is being used, consider whether the unexpected distribution results might be the result of how well the server is meeting the goals of its service class, and the amount of workload available on the system given the importance of its service class. See the z/OS Communications Server: IP Configuration Guide for an overview of how WLM determines server recommendations and how they are used by TCP/IP. For more detailed information about sysplex routing services, see z/OS MVS Planning: Workload Management.
   - If some entries have a low TSR value, consider whether network or server performance problems might be affecting distribution. Examine the TCSR, CER, and SEF values in the DETAIL output for these entries.
     – A low TCSR value indicates a connectivity problem between the sysplex distributor stack and the target stack for those particular DVIPA, Port,

and Destination entries. To diagnose, verify connectivity across the dynamic XCF interface (HiperSockets, IUTSAMEH, XCF interfaces) or that VIPAROUTE is working properly by interface-specific ping.

- If the SEF value is low but the CER is not, then the application on this target is having problems with accepting new connections. Take the following steps to diagnose:
  a. Verify that the application is accepting connections by connecting to another IP address on the target.
  b. If the connection is not accepted, verify that the application is getting dispatched correctly.
  c. If the connection is accepted, verify that the application can properly handle the load (amount of work, connections, tasks).
- If the SEF and CER values are low, then the target stack is having problems with establishing connections with one or more clients. Take the following steps to diagnose:
  a. Find the client to which the connection is failing. Because **netstat conn** does not show connections in `syn sent`, the best way to find the clients that are not completing the three-way handshake is a packet trace (other means of finding the client IP are possible).
  b. If the client is known, verify that the route from the target to the client is correct by reviewing the target stack routing table. If the route from the target to the client is correct, verify that traceroute shows the correct path is being followed to the client.
  c. If the route from the target to the client is correct, verify that traceroute shows the correct path is being followed to the client.
- If all entries representing distribution to the same target are very low, or 0, this might indicate that the target stack is experiencing problems.
- If you used a VIPAROUTE definition to specify the route from the distributor to the target, check the specified route to verify that it is active.

- If SERVERWLM is being used as the distribution method and a server has a WLM weight of 0, verify that the server is using the appropriate WLM Policy and that the system is not too overloaded to enable the server to meet its policy goals. See z/OS Communications Server: IP Configuration Guide for an overview of how WLM determines server recommendations and how they are used by TCP/IP. For more detailed information about sysplex routing services, see z/OS MVS Planning: Workload Management.
- If the unexpected distribution results have not yet been explained and Sysplex Distributor Performance Policies have been defined using Policy Agent, consider whether the distribution might be caused by two network performance issues (TCP retransmissions and timeouts).
- If Sysplex Distributor Routing Policies have been defined using Policy Agent, consider whether the definition of that policy is affecting the connection distribution. After determining which connections are not being distributed correctly, run D TCPIP,TCPCS,NET,VCRT,DETAIL (see step 10 on page 395) to determine the policy action to which each connection maps. Look at the QoS weights for those policy actions in the VDPT DETAIL display to see whether they are unusually low. The Policy Agent log on the target stack can display for each DVIPA/Port the QoS service level fractions used to modify the QoS weight. It can also display the calculations that caused a QoS fraction to be set abnormally high (such as connection limit exceeded or throughput exceeded). See "Diagnosing Policy Agent problems" on page 661 for more information.

10. Run the Netstat VCRT/-V command on the distributing stack to check whether there are any active connections that are being routed by the distributor. If you run the command with the keyword DETAIL (**d tcpip,tcpcs,net,vcrt,detail**) you can see the policy rule and policy action that each connection maps to.

    If the VCRT table shown in Figure 38 is empty, then connection requests might not be reaching the distributor. Check for a routing problem from the client to the distributor.

    If you see expected entries in the table, note the dynamic XCF address and proceed to step 11. Figure 38 shows an example:

```
D TCPIP,TCPCS,NET,VCRT,DETAIL
EZD0101I NETSTAT CS V1R7 TCPCS 758
DYNAMIC VIPA CONNECTION ROUTING TABLE:
DEST:      201.2.10.11..21
  SOURCE:  203.110.1.1..1031
  DESTXCF: 193.15.1.1
    POLICYRULE:    *NONE*
    POLICYACTION:  *NONE*
DEST:      201.2.10.12..21
  SOURCE:  203.110.1.1..1033
  DESTXCF: 193.9.200.2
DEST:      201.2.10.13..5000
  SOURCE:  203.110.1.1..0
  DESTXCF: 193.15.1.1
    CFGTIMAFF: 0045  TIMAFFCNT: 0000000002  TIMAFFLFT: 0000
DEST:      201.2.10.13..5000
  SOURCE:  203.110.1.1..1029
  DESTXCF: 193.15.1.1
    POLICYRULE:    *NONE*
    POLICYACTION:  *NONE*
DEST:      201.2.10.13..5000
  SOURCE:  203.110.1.1..1030
  DESTXCF: 193.15.1.1
    POLICYRULE:    *NONE*
    POLICYACTION:  *NONE*
```

Figure 38. d tcpip,tcpcs,net,vcrt,detail example

11. Go to the target stacks represented by the DESTXCF ADDR field in the VCRT or VDPT display and run the Netstat ALLCONN(/-a),IPA=201.2.10.12 display command to see the connections on the target stack. Figure 39 shows an example:

```
D TCPIP,TCPCS2,NET,ALLCONN,IPA=201.2.10.12
EZD0101I NETSTAT CS V1R7 TCPCS2 846
USER ID  CONN     STATE
FTPD1    000000F3 ESTBLSH
  LOCAL SOCKET:   ::FFFF:201.2.10.12..21
  FOREIGN SOCKET: ::FFFF:203.110.1.1..1033
1 OF 1 RECORDS DISPLAYED
```

Figure 39. Netstat ALLConn/-a example with IPAddr/-I filter value of 201.2.1.12

Tip: For a variety of reasons, the VCRT and ALLCONN displays might not match exactly. For example, with short-lived connections such as Web connections, an entry might show up in one display but be gone by the time the second display is run. Also, the distributing stack places an entry into the Dynamic VIPA Connection Routing Table when it first forwards a connection request. A busy server might reject these connection requests, and therefore cause a temporary mismatch in the two displays.

Steps for diagnosing problems using DVIPAs in source IP address selection for TCPconnections problems

Investigating problems that are related to which source IP address is chosen for outbound TCP connections depends on which options you configured.

Procedure

1. If you are using the TCPSTACKSRCVIPA function, then run the Netstat CONFIG/-f command on the stack in question to verify that the sysplex-wide dynamic source VIPA was configured as expected. In other words, verify that IPCONFIG/IPCONFIG6 SOURCEVIPA is set to YES and that IPCONFIG/IPCONFIG6 TCPSTACKSRCVIPA is specified with the correct address or interface. If you are using TCPSTACKSOURCEVIPA with a distributed DVIPA, run the Netstat CONFIG/-f command on the distributor stack and on the target stacks. Figure 40 shows an example:

```
IP CONFIGURATION TABLE:
FORWARDING: YES    TIMETOLIVE: 00064  RSMTIMEOUT:  00060
ARPTIMEOUT: 01200  MAXRSMSIZE: 65535  FORMAT:      LONG
IGREDIRECT: YES    SYSPLXROUT: YES    DOUBLENOP:   NO
STOPCLAWER: NO     SOURCEVIPA: YES
MULTIPATH:  NO     PATHMTUDSC: NO     DEVRTRYDUR:  0000000090
DYNAMICXCF: YES
  IPADDR/PREFIXLEN: 193.15.1.1/24       METRIC: 02
IQDIOROUTE: NO
TCPSTACKSRCVIPA: 203.15.1.1
IPV6 CONFIGURATION TABLE:
FORWARDING:    YES  HOPLIMIT:   00255  IGREDIRECT:  YES
SOURCEVIPA:    YES  MULTIPATH:  NO     ICMPERRLIM:  00003
IGRTRHOPLIMIT: NO
DYNAMICXCF: YES
   IPADDR: 2001:0DB8::151:0
   INTFID: 0006:0007:0008:0009
TCPSTACKSRCVIPA: DVIPA1
```

Figure 40. Netstat CONFIG/-f example

2. If you are using the SRCIP function to specify source IP addressing for specified jobnames or destinations, then run the Netstat SRCIP/-J command to display the SRCIP configuration. Verify that either the jobname for the application performing the outbound CONNECT() or the destination address for the CONNECT() matches an entry in the SRCIP configuration. See z/OS Communications Server: IP Configuration Reference for the order of precedence that is followed if an outbound connection matches more than one entry in the SRCIP configuration.

   If you have configured distributed DVIPAs on SRCIP rules and outbound connections are failing with EADDRNOTAVAIL and JRSRCIPDistDVIPA , EXPLICITBINDPORTRANGE processing is either not configured or not working properly. See “Steps for diagnosing problems with EXPLICITBINDPORTRANGE processing” on page 401 for more information.

```
D TCPIP,,N,SRCIP
EZD0101I NETSTAT CS V1R8 TCPCS 745
SOURCE IP ADDRESS BASED ON JOB NAME:
JOB NAME  TYPE  SOURCE
--------  ----  ------
USER*     IPV4  203.15.2.1
USER*     IPV6  2003::15:1:1

SOURCE IP ADDRESS BASED ON DESTINATION:
DESTINATION: 192.1.1.98
  SOURCE:    203.15.2.2
DESTINATION: 2001:DB8:10::82:2:2
  SOURCE:    2003::15:1:2
4 OF 4 RECORDS DISPLAYED
END OF THE REPORT
```

Figure 41. Netstat SRCIP/-J example

3. Create an outbound connection. Use the Netstat ALLConn/-a command to confirm that the correct source IP address was used.

   Tip: TCPSTACKSOURCEVIPA and SRCIP specifications can be overridden. For example, a match to an SRCIP entry will override a TCPSTACKSOURCEVIPA specification. If your TCPSTACKSOURCEVIPA or SRCIP configuration is correct but you are not getting the expected source IP address, see the source IP address selection information in z/OS Communications Server: IP Configuration Guide for the hierarchy of the various ways that the source IP address of an outbound packet is determined.

Steps for diagnosing problems with the SYSPLEX-wide ephemeral port assignment for distributed DVIPAs

Diagnose problems with the SYSPLEXPORTS setting.

Procedure

Perform the following steps:

1. Run the Netstat VIPADCFG/-F command on the distributor stack to confirm that SYSPLEXPORTS was specified for all distributed DVIPAs as expected.

```
D TCPIP,TCPCS,NET,VIPADCFG
EZD0101I NETSTAT CS V1R7 TCPCS 862
DYNAMIC VIPA INFORMATION:
  VIPA DEFINE:
    IPADDR/PREFIXLEN: 203.15.1.1/24
      MOVEABLE: IMMEDIATE  SRVMGR: NO
    IPADDR/PREFIXLEN: 203.15.1.2/24
      MOVEABLE: IMMEDIATE  SRVMGR: NO
    INTFNAME: DVIPA1
      IPADDR: 2001:0DB8:1::1
        MOVEABLE: IMMEDIATE  SRVMGR: N/A
  VIPA DISTRIBUTE:
    DEST:        203.15.1.1..4011
      DESTXCF:   ALL
        SYSPT:   YES  TIMAFF: NO    FLG: BASEWLM
    DEST:        203.15.1.2..245
      DESTXCF:   ALL
        SYSPT:   NO   TIMAFF: NO    FLG: BASEWLM
    DESTINTF:    DVIPA1
      DEST:      2001:0DB8:1::1
        DESTXCF: ALL
          SYSPT: YES  TIMAFF: NO    FLG: BASEWLM
```

Figure 42. Diagnosing SYSPLEXPORTS problems

In the preceding display, the distributed DVIPAs 203.15.1.1 and 2001:0DB8:1::1 were enabled with SYSPLEXPORTS(SYSPT is Yes), while 203.15.1.2 was not (SYSPT is NO).

2. Verify from the system log that the following message was issued:

   ```
   IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBEPORT
   ```

   If subplexing is being used, the message will be:

   ```
   IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBEPORTvvtt
   ```

   where *vv* is the VTAM subplex group ID and *tt* is the TCP/IP subplex group ID. If no VTAM subplex group ID was specified at VTAM startup, but a TCP/IP subplex group ID was specified on the GLOBALCONFIG statement in the TCP/IP Profile, then the structure name is EZBEPORT01*tt*. If a VTAM subplex group ID was specified, but no TCP/IP subplex group ID was specified, then the structure name is EZBEPORT*vv*.

   If this message was not issued and Netstat VIPADCFG/-F shows that the SYSPLEXPORTS field was specified, see z/OS Communications Server: SNA Network Implementation Guide for more information about defining EZBEPORTvvtt with the coupling facility.

3. Bind to an ephemeral port and then create an outbound connection with the source IP address of the SYSPLEXPORTS distributed DVIPA. Do the following to verify SYPLEXPORTS is working correctly:

   a. Issue the Netstat ALLConn/-a command to verify the connection on the target stack.

   b. Issue the Netstat VCRT/-V command to confirm that the distributing stack is aware of the connection.

   c. Issue the VTAM DISPLAY NET,STATS command, specifying the full name of the EZBEPORT structure to confirm that the coupling facility is managing ports for this distributed DVIPA. For ephemeral ports, the coupling facility assigns a block of 64 ports to the TCP/IP stack. For example:

```
D NET,STATS,TYPE=CFS,STRNAME=EZBEPORT1121,LIST=ALL,SCOPE=ALL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = STATS,TYPE=CFS 180
IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBEPORT1121
IST1797I STRUCTURE TYPE = LIST
IST1517I LIST HEADERS = 1024 - LOCK HEADERS = 1024
IST1373I STORAGE ELEMENT SIZE = 256
IST924I -------------------------------------------------------------
IST1374I                            CURRENT     MAXIMUM  PERCENT
IST1375I STRUCTURE SIZE               8192K      15104K     *NA*
IST1376I STORAGE ELEMENTS               128       22400        0
IST1377I LIST ENTRIES                     5         700        0
IST924I -------------------------------------------------------------
ISTREW1I EXPLICITBINDPORTRANGE - START: 50000 END: 50255
IST1823I LIST DVIPA SYSNAME  TCPNAME               # ASSIGNED PORTS
IST1824I    0 EXPLICITBINDPORTRANGE                              64
IST1825I            VIC015   TCPCS                               64
IST1826I               PORTS: 50000 50001 50002 50003 50004 50005
IST1827I                      50006 50007 50008 50009 50010 50011
IST1827I                      50012 50013 50014 50015 50016 50017
IST1827I                      50018 50019 50020 50021 50022 50023
IST1827I                      50024 50025 50026 50027 50028 50029
IST1827I                      50030 50031 50032 50033 50034 50035
IST1827I                      50036 50037 50038 50039 50040 50041
IST1827I                      50042 50043 50044 50045 50046 50047
IST1827I                      50048 50049 50050 50051 50052 50053
IST1827I                      50054 50055 50056 50057 50058 50059
IST1827I                      50060 50061 50062 50063
IST1824I    1 203.15.1.1                                         64
IST1825I            VIC015   TCPCS                               64
IST1826I               PORTS:  1024  1025  1026  1027  1028  1029
IST1827I                       1030  1031  1032  1033  1034  1035
IST1827I                       1036  1037  1038  1039  1040  1041
IST1827I                       1042  1043  1044  1045  1046  1047
IST1827I                       1048  1049  1050  1051  1052  1053
IST1827I                       1054  1055  1056  1057  1058  1059
IST1827I                       1060  1061  1062  1063  1064  1065
IST1827I                       1066  1067  1068  1069  1070  1071
IST1827I                       1072  1073  1074  1075  1076  1077
IST1827I                       1078  1079  1080  1081  1082  1083
IST1827I                       1084  1085  1086  1087
IST1824I    2 2001:0DB8:1::1                                     64
IST1825I            VIC015   TCPCS                               64
IST1826I               PORTS:  1024  1025  1026  1027  1028  1029
IST1827I                       1030  1031  1032  1033  1034  1035
IST1827I                       1036  1037  1038  1039  1040  1041
IST1827I                       1042  1043  1044  1045  1046  1047
IST1827I                       1048  1049  1050  1051  1052  1053
IST1827I                       1054  1055  1056  1057  1058  1059
IST1827I                       1060  1061  1062  1063  1064  1065
IST1827I                       1066  1067  1068  1069  1070  1071
IST1827I                       1072  1073  1074  1075  1076  1077
IST1827I                       1078  1079  1080  1081  1082  1083
IST1827I                       1084  1085  1086  1087
IST314I END
```

Figure 43. VTAM NET,STATS example

Diagnosing problems with the SYSPLEX-wide ephemeral port assignment for EXPLICITBINDPORTRANGE processing

This topic contains two subtopics:

- "Steps for determining an optimal range for the EXPLICITBINDPORTRANGE parameter" on page 400

- "Steps for diagnosing problems with EXPLICITBINDPORTRANGE processing" on page 401

Steps for determining an optimal range for the EXPLICITBINDPORTRANGE parameter

Determine an optimal range for the EXPLICITBINDPORTRANGE parameter.

Procedure

Perform the following steps:

1. Change your existing configuration to specify the GLOBALCONFIG EXPLICITBINDPORTRANGE parameter on all stacks that you anticipate will be participating in explicit bind port range processing. As a guideline, the port range size should be at least large enough to allow for 2 blocks of ports to be in use by each participating TCP/IP stack (128 * number of TCP/IP stacks using the explicit bind port range). If you are using an SRCIP block, do not initially make the change to use distributed DVIPAs on the DESTINATION rules.
2. Start all of the stacks, servers and clients to reach the typical steady state connection load for your sysplex environment. All connections from sockets that were explicitly bound to the IPv4 address, INADDR_ANY, or the IPv6 unspecified address (in6addr_any) and port 0 will use ports from the new range.
3. Periodically check to determine how many ports from this new range are in use by issuing a D NET,STATS,TYPE=CFS,STRNAME=EZBEPORTvvtt,LIST=0 command (See Figure 44).

```
D NET,STATS,TYPE=CFS,STRNAME=EZBEPORT1121,LIST=0
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = STATS,TYPE=CFS 180
IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBEPORT1121
IST1797I STRUCTURE TYPE = LIST
IST1517I LIST HEADERS = 1024 - LOCK HEADERS = 1024
IST1373I STORAGE ELEMENT SIZE = 256
IST924I -------------------------------------------------------------
IST1374I                            CURRENT     MAXIMUM  PERCENT
IST1375I STRUCTURE SIZE               8192K      15104K     *NA*
IST1376I STORAGE ELEMENTS              128       22400       0
IST1377I LIST ENTRIES                    5         700       0
IST924I -------------------------------------------------------------
ISTREW1I EXPLICITBINDPORTRANGE - START: 50000 END: 50255
IST1823I LIST DVIPA SYSNAME  TCPNAME               # ASSIGNED PORTS
IST1824I    0 EXPLICITBINDPORTRANGE                             192
IST1825I            VIC015   TCPCS                               64
IST1825I            VIC021   TCPCS11                            128
IST314I END
```

Figure 44. VTAM NET,STATS,LIST=0 example

4. Check for message EZD1296 which will be issued if local ephemeral ports were used for connections because no explicit bind ports were available from the active EXPLICITBINDPORTRANGE parameter.
5. If message EZD1296 is not issued and if the number of allocated ports is consistently less than the total port range, then proceed to the next step, otherwise:
   a. Change the GLOBALCONFIG EXPLICITBINDPORTRANGE parameter to use a larger explicit bind port range; as a guideline increase the size by at least 64 for each participating stack (64 multiplied by the number of TCP/IP stacks)

b. Issue a VARY TCPIP,,OBEYFILE command to change the range on each of the stacks and return to step 3 above.

6. Change your SRCIP block to use distributed DVIPAs on DESTINATION rules.

Steps for diagnosing problems with EXPLICITBINDPORTRANGE processing

If you configured distributed DVIPAs on SRCIP rules and outbound connections are failing with EADDRNOTAVAIL and JRSRCIPDistDVIPA, EXPLICITBINDPORTRANGE processing is not working.

About this task

There are several possible reasons for this problem:

- EXPLICITBINDPORTRANGE parameter is not configured on the stack.
- EXPLICITBINDPORTRANGE parameter is configured but:
  - The stack did not connect to the EZBEPORT structure in the coupling facility
  - The stack has lost access to the EZBEPORT structure
  - The stack is running in a CINET environment with more than 1 stack and stack affinity was not established
  - There are no available ports in the range (the range is exhausted)
  - The application is bound explicitly to an ephemeral port (equal to or greater than the low port value that is specified on the EPHEMERALPORTS parameter of the TCPCONFIG profile statement, or 1024 by default) that is not reserved for this job by the PORT or PORTRANGE profile statement

Procedure

Use the following steps to determine and correct the problem:

1. Issue the D TCPIP,tcp_stackname,SYSPLEX,PORTS command to determine the configured EXPLICITBINDPORTRANGE value for this stack and the active EXPLICITBINDPORTRANGE value in the sysplex (or subplex). If the command response indicates "No EXPLICITBINDPORTRANGE is configured on this stack", see z/OS Communications Server: IP Configuration Reference, under the section for the GLOBALCONFIG statement, for information about enabling the EXPLICITBINDPORTRANGE.

   Tip: If the active port range does not match the configured port range for this stack, it means that another stack that was started after this stack had a different range defined in the GLOBALCONFIG EXPLICITBINDPORTRANGE parameter, or a VARY OBEYFILE command was processed on another stack that specified a GLOBALCONFIG EXPLICITBINDPORTRANGE parameter with a different range. You should try to ensure that all stacks participating in EXPLICITBINDPORTRANGE parameter processing specify the same port range. This can be done by specifying the GLOBALCONFIG EXPLICITBINDPORTRANGE statement in a file that is included in each stack's TCP/IP profile using an INCLUDE statement.

2. Verify from the system log that the following message was issued:
   - IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBEPORT
   - If subplexing is being used, the message will be:

     ```
     IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBEPORTvvtt
     ```

     The vv value is the VTAM subplex group ID and the tt value is the TCP/IP subplex group ID. If no VTAM subplex group ID was specified at VTAM

startup, but a TCP/IP subplex group ID was specified on the GLOBALCONFIG statement in the TCP/IP profile, then the structure name is EZBEPORT01tt. If a VTAM subplex group ID was specified, but no TCP/IP subplex group ID was specified, then the structure name is EZBEPORTvv.

- If this message was not issued and D TCPIP,,SYSPLEX,PORTS shows that an EXPLICITBINDPORTRANGE parameter was configured, see z/OS Communications Server: SNA Network Implementation Guide for more information about defining EZBEPORTvvtt with the coupling facility.
- If the IST1370I message was issued (the stack did connect to the EZBEPORT structure), and D TCPIP,,SYSPLEX,PORTS shows that an EXPLICITBINDPORTRANGE parameter was configured but that no active EXPLICITBINDPORTRANGE is available from this stack, the stack may have lost connectivity to the EZBEPORT structure either as a result of a structure rebuild or a structure disconnect. In the console log, check for any failure or rebuild messages referencing the EZBEPORT structure. If a structure rebuild was in process for the EZBEPORT structure in use by this stack, wait for the rebuild to complete. If VTAM lost connectivity to the structure, issue the VARY NET,CFS,ACTION=CONNECT,STRNAME=structure_name command to reestablish connectivity to the structure

3. Bind to the IPv4 address, INADDR_ANY, or the IPv6 unspecified address (in6addr_any) and port 0. Issue the D NET,STATS,TYPE=CFS,STRNAME=EZBEPORTvvtt,LIST=0 command to confirm that the coupling facility has ports in the EXPLICITBINDPORTRANGE parameter allocated for the stack on which you issued the bind, and issue the Netstat ALLConn/-a command to determine whether the port assigned to your application for this bind was one in the explicit bind port range.

```
netstat -a -p tcpcs1
MVS TCP/IP NETSTAT CS V1R9       TCPIP Name: TCPCS1          13:21:04
User Id  Conn     State
-------  ----     -----
BPXOINIT 00000017 Listen
  Local Socket:   0.0.0.0..10007
  Foreign Socket: 0.0.0.0..0
USER11   0000002B Closed
  Local Socket:   0.0.0.0..50001
  Foreign Socket: 0.0.0.0..0
USER11   00000021 Closed
  Local Socket:   0.0.0.0..50000
  Foreign Socket: 0.0.0.0..0
SYSLOGD8 00000018 UDP
  Local Socket:   ::..514
  Foreign Socket: *..*
```

   If a port from the explicit bind port range was not allocated, check if you are running in a CINET environment in which more than one TCP/IP stack is being managed by CINET and stack affinity has not been established. Explicit bind port range processing is not supported in such a configuration.

4. Check the system log for message "EZD1296I EXPLICITBINDPORTRANGE exhausted" which indicates that the number of ports in the EXPLICITBINPORTRANGE parameter is not large enough. The coupling facility was unable to allocate a port from this range and a stack ephemeral port was allocated instead. This may be a temporary situation because EBPR ports are eventually returned to the coupling facility after sockets bound to them are closed.

   Tip: Message EZD1296I is not issued more than once every 5 minutes. If this message is issued multiple times, you should consider enlarging the number of

ports for the EXPLICITBINDPORTRANGE parameter. See “Steps for determining an optimal range for the EXPLICITBINDPORTRANGE parameter” on page 400 for more information.

5. Issue the Netstat ALLCONN/-a command to display the local socket IP address and port the application is bound to, and Netstat PORTLIST/-o command to display the ports that are reserved. If you want to add the applications local port to the list of reserved ports, use the PORT or PORTRANGE profile statement.

Diagnosing sysplex-wide security association (SWSA) problems

This topic describes methods for diagnosing SWSA problems.

Steps for diagnosing sysplex-wide security association (SWSA) problems

Use the following information to aid with diagnosing Sysplex-wide Security Association (SWSA) specifically.

Before you begin

Ensure that you have consistent IPSec policies on all participating systems, which include the following:

- Distributing stacks, target stacks and backup stacks.
- Certificates identifying hosts must be available on all distributing and backup hosts. This is most easily accomplished by sharing the SAF certificate repository between the processors in the sysplex.

See z/OS Communications Server: IP Configuration Guide for information about configuring IP security policy on an IPSECURITY stack.

Procedure

Perform the following steps to diagnose SWSA problems.

1. Code the DVIPSEC option on the owning and backup stacks to take advantage of SWSA. Do the following on the owning and backup stacks:
   - Use the **netstat,config** command to confirm that IPSECURITY was specified on the IPCONFIG statement and, if appropriate, on the IPCONFIG6 statement.

     Figure XX. netstat,config example

     ```
     D TCPIP,,NETSTAT,CONFIG
     NETSTAT CONFIG
     MVS TCP/IP NETSTAT CS V2R1        TCPIP Name: TCPCS
     ...
     IP Configuration Table:
     Forwarding: Yes    TimeToLive: 00064  RsmTimeOut:  00060
     IPSecurity: Yes
     ...
     IPv6 Configuration Table:
     Forwarding:    Yes  HopLimit:   00255  IgRedirect:  No
     SourceVipa:    Yes  MultiPath:  Conn   IcmperrLim:  00003
     IgRtrHopLimit: No
     IpSecurity: Yes
     ```

   - Use the **ipsec -f** command to confirm that DVIPSEC was specified on the IPSEC statement.

```
# ipsec -f disp
CS V2R1 ipsec  Stack Name: TCPCS1  Fri May 27 10:48:47 2011
Primary:  Filter          Function: Display           Format:   Detail
Source:   Stack Profile   Scope:    Current           TotAvail: 2
Logging:  Off             Predecap: Off               DVIPSec:  Yes
```

Figure 45. ipsec -f example

2. Verify from the system log for the distributing and target stacks (for sysplex distribution of IPSec workload) and the primary and backup stacks (for dynamic tunnel recovery) that an IST1370I message like the following was issued:

   ```
   IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBDVIPA
   ```

   If subplexing is being used, the message is:

   ```
   IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBDVIPAvvtt
   ```

   where *vv* is the VTAM subplex group ID and *tt* is the TCP/IP subplex group ID. If no VTAM subplex group ID was specified at VTAM startup, but a TCP/IP subplex group ID was specified on the GLOBALCONFIG statement in the TCP/IP Profile, then the structure name is EZBDVIPA01*tt*. If a VTAM subplex group ID was specified, but no TCP/IP subplex group ID was specified, then the structure name is EZBDVIPA*vv*.

   For SWSA functions to work correctly, the stacks involved must be connected to the EZBDVIPA*vvtt* coupling facility structure. If this message was not issued, see z/OS Communications Server: SNA Network Implementation Guide for information about setting up the sysplex environment for VTAM function and defining EZBDVIPAvvtt with the coupling facility.

3. For sysplex distribution of IPSec traffic, the target stacks must have a copy of the dynamic tunnel, called a shadow tunnel, that matches the dynamic tunnel on the distributing stack. Do the following:
   a. Use the following command to verify that a dynamic tunnel is active on distributing stacks:

```
# ipsec -y display
CS V2R1 ipsec  Stack Name: TCPCS1  Fri May 27 13:22:15 2011
Primary:  Dynamic tunnel  Function: Display           Format:   Detail
Source:   Stack           Scope:    Current           TotAvail: 2

TunnelID:                     Y6
Generation:                   1
IKEVersion:                   2.0
ParentIKETunnelID:            K5
VpnActionName:                VPN-14-Transport
LocalDynVpnRule:              n/a
State:                        Active
HowToEncap:                   Transport
LocalEndPoint:                10.93.1.8
RemoteEndPoint:               10.83.4.1
LocalAddressBase:             10.93.1.8
LocalAddressPrefix:           n/a
LocalAddressRange:            n/a
RemoteAddressBase:            10.83.4.1
RemoteAddressPrefix:          n/a
RemoteAddressRange:           n/a
HowToAuth:                    ESP
 AuthAlgorithm:               HMAC-SHA1
 AuthInboundSpi:              1337238202 (0x4FB4A2BA)
 AuthOutboundSpi:             3188885928 (0xBE1289A8)
HowToEncrypt:                 3DES-CBC
 KeyLength:                   n/a
 EncryptInboundSpi:           1337238202 (0x4FB4A2BA)
 EncryptOutboundSpi:          3188885928 (0xBE1289A8)
Protocol:                     ALL(0)
LocalPort:                    n/a
LocalPortRange:               n/a
RemotePort:                   n/a
RemotePortRange:              n/a
Type:                         n/a
TypeRange:                    n/a
Code:                         n/a
CodeRange:                    n/a
OutboundPackets:              0
OutboundBytes:                0
InboundPackets:               0
InboundBytes:                 0
Lifesize:                     0K
LifesizeRefresh:              0K
CurrentByteCount:             0b
LifetimeRefresh:              2011/05/27 16:40:23
LifetimeExpires:              2011/05/27 16:58:50
CurrentTime:                  2011/05/27 12:58:56
VPNLifeExpires:               2011/05/28 12:58:50
NAT Traversal Topology:
  UdpEncapMode:               No
  LclNATDetected:             No
  RmtNATDetected:             No
  RmtNAPTDetected:            No
  RmtIsGw:                    n/a
  RmtIsZOS:                   n/a
  zOSCanInitP2SA:             n/a
  RmtUdpEncapPort:            n/a
  SrcNATOARcvd:               n/a
  DstNATOARcvd:               n/a
PassthroughDF:                n/a
PassthroughDSCP:              n/a
*********************************************************************
```

Figure 46. ipsec -y example

b. Use the following command to verify that a shadow tunnel is active on target stacks:

```
# ipsec -y display -s
CS V2R1 ipsec  Stack Name: TCPCS2  Fri May 27 13:22:23 2011
Primary:  Dynamic tunnel  Function: display (shadows)  Format:   Detail
Source:   Stack           Scope:    Current            TotAvail: 2

TunnelID:                     Y6
Generation:                   1
IKEVersion:                   2.0
ParentIKETunnelID:            K5
VpnActionName:                VPN-14-Transport
LocalDynVpnRule:              n/a
State:                        Active
HowToEncap:                   Transport
LocalEndPoint:                10.93.1.8
RemoteEndPoint:               10.83.4.1
LocalAddressBase:             10.93.1.8
LocalAddressPrefix:           n/a
LocalAddressRange:            n/a
RemoteAddressBase:            10.83.4.1
RemoteAddressPrefix:          n/a
RemoteAddressRange:           n/a
HowToAuth:                    ESP
 AuthAlgorithm:               HMAC-SHA1
 AuthInboundSpi:              1337238202 (0x4FB4A2BA)
 AuthOutboundSpi:             3188885928 (0xBE1289A8)
HowToEncrypt:                 3DES-CBC
 KeyLength:                   n/a
 EncryptInboundSpi:           1337238202 (0x4FB4A2BA)
 EncryptOutboundSpi:          3188885928 (0xBE1289A8)
Protocol:                     ALL(0)
LocalPort:                    n/a
LocalPortRange:               n/a
RemotePort:                   n/a
RemotePortRange:              n/a
Type:                         n/a
TypeRange:                    n/a
Code:                         n/a
CodeRange:                    n/a
OutboundPackets:              0
OutboundBytes:                0
InboundPackets:               0
InboundBytes:                 0
Lifesize:                     0K
LifesizeRefresh:              0K
CurrentByteCount:             0b
LifetimeRefresh:              2011/05/27 16:40:23
LifetimeExpires:              2011/05/27 16:58:50
CurrentTime:                  2011/05/27 13:22:24
VPNLifeExpires:               2011/05/28 12:58:50
NAT Traversal Topology:
  UdpEncapMode:               No
  LclNATDetected:             No
  RmtNATDetected:             No
  RmtNAPTDetected:            No
  RmtIsGw:                    n/a
  RmtIsZOS:                   n/a
  zOSCanInitP2SA:             n/a
  RmtUdpEncapPort:            n/a
  SrcNATOARcvd:               n/a
  DstNATOARcvd:               n/a
PassthroughDF:                n/a
PassthroughDSCP:              n/a
```

Figure 47. ipsec -y display -s example

4. To confirm that the coupling facility has the information about the tunnels in the event a recovery is necessary, use the following VTAM command, specifying the full name of the EZBDVIPA structure:

   ```
   d net,stats,type=cfs,strname=ezbdvipa1121,dvipa=10.93.1.8
   ```

   The following output is displayed:

   ```
   IST097I DISPLAY ACCEPTED
   IST350I DISPLAY TYPE = STATS,TYPE=CFS
   IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBDVIPA1121
   IST1797I STRUCTURE TYPE = LIST
   IST1517I LIST HEADERS = 2048 - LOCK HEADERS = 0
   IST1373I STORAGE ELEMENT SIZE = 256
   IST924I -------------------------------------------------------------
   IST1374I                             CURRENT     MAXIMUM  PERCENT
   IST1375I STRUCTURE SIZE               15104K      50176K    *NA*
   IST1376I STORAGE ELEMENTS                 48       28592       0
   IST1377I LIST ENTRIES                      6        2902       0
   IST924I -------------------------------------------------------------
   IST1834I LIST DVIPA SYSNAME  TCPNAME    #ENTRIES    TGCOUNT SEQNUMBER
   IST1835I    1 10.93.1.8
   IST1837I            MVS187   TCPCS1            1                   1
   IST1835I    2 10.93.1.8
   IST1836I            MVS187   TCPCS1            1          0
   IST314I END
   ```

   Information about the dynamic tunnels that are used in SWSA is kept in the coupling facility structure in the event that a recovery of the tunnel is necessary. For example, the recovery information is used when a DVIPA is taken over by another stack in the sysplex.

   For more information, see DISPLAY STATS in z/OS Communications Server: SNA Operation.

   For IPSec connections to continue functioning with that DVIPA, the tunnel has to be recovered by the same stack that took over the dynamic VIPA.

   The list entry for the DVIPA (list 2 above) shows the system and stack for which the coupling facility is maintaining information about the tunnel.

5. Use the following VTAM command, specifying the full name of the EZBDVIPA structure, to confirm that the coupling facility is managing the replay count:

   ```
   d net,stats,type=cfs,strname=ezbdvipa1121,list=all
   ```

   For sysplex distribution of IPSec traffic, the dynamic tunnel replay count (sequence number) is maintained in the EZBDVIPAvvtt coupling facility structure. The distributing stack dynamic tunnel and all the target stack shadow tunnels share the replay count.

   The following output is displayed:

   ```
   D NET,STATS,TYPE=CFS,STRNAME=EZBDVIPA1121,LIST=ALL

   IST097I DISPLAY ACCEPTED
   IST350I DISPLAY TYPE = STATS,TYPE=CFS
   IST1370I NETA.SSCP1A IS CONNECTED TO STRUCTURE EZBDVIPA1121
   IST1797I STRUCTURE TYPE = LIST
   IST1517I LIST HEADERS = 2048 - LOCK HEADERS = 0
   IST1373I STORAGE ELEMENT SIZE = 256
   IST924I -------------------------------------------------------------
   IST1374I                             CURRENT     MAXIMUM  PERCENT
   IST1375I STRUCTURE SIZE               15104K      50176K    *NA*
   IST1376I STORAGE ELEMENTS                 48       28592       0
   IST1377I LIST ENTRIES                      6        2902       0
   IST924I -------------------------------------------------------------
   IST1834I LIST DVIPA SYSNAME  TCPNAME    #ENTRIES    TGCOUNT SEQNUMBER
   IST1835I    1 10.93.1.8
   ```

```
IST1837I            MVS187   TCPCS1            1                    1
IST1835I    2 10.93.1.8
IST1836I            MVS187   TCPCS1            1          0
IST314I END
```

The list entry for the dynamic VIPA with a value in the SEQNUMBER column confirms that this tunnel's replay count is managed by the coupling facility.

Steps for diagnosing sysplex routing problems

Determine why packets are not distributed properly in your sysplex.

Procedure

Perform the following steps:

1. Run the Netstat VIPADyn VIPAROUTE/-v VIPAROUTE command on the distributing stack to see what type of route is used for distributing packets to target stacks.

```
D TCPIP,TCPCS,NET,VIPADYN,VIPAROUTE
EZD0101I NETSTAT CS V1R7 TCPCS
VIPA ROUTE:
  DESTXCF: 193.1.3.94
    TARGETIP: 213.5.1.1
    RTSTATUS: ACTIVE
  DESTXCF: 193.1.4.94
    TARGETIP: 213.6.2.2
    RTSTATUS: INACTIVE
  DESTXCF: 2EC0::943:F003
    TARGETIP: 1EC0::5:1:1
    RTSTATUS: ACTIVE
  DESTXCF: 2EC0::943:F004
    TARGETIP: 1EC0::6:2:2
    RTSTATUS: INACTIVE
4 OF 4 RECORDS DISPLAYED
```

Figure 48. Netstat VIPADyn example

   - If there is no VIPA ROUTE entry, IP packets that are distributed by Sysplex Distributor to target stacks use dynamic XCF interfaces. Use the Netstat ROUTe/-r command on the distributing stack to see other routing failure problems.
   - If there is a VIPA ROUTE entry defined for a target stack and the RtStatus field shows `Active`, IP packets that are distributed by Sysplex Distributor to that target stack use the normal IP routing tables to determine the best available route.
   - If there is a VIPA ROUTE entry defined for that target stack and the RtStatus field shows `Unavail`, the defined target IP address in the route entry is not available yet. This could be because the target stack is currently active, but the target IP address is not defined in that target stack. All packets to that target stack use dynamic XCF interfaces. This is likely to be a configuration error that should be investigated. EZD1173I is issued when the stack detects this problem.
   - If there is a VIPA ROUTE entry defined for a target stack and the RtStatus field shows `Inactive`, no route exists to that target stack. See z/OS Communications Server: IP System Administrator's Commands for more information about the RtStatus field.

2. Run the Netstat ROUTe/-r command on the distributing stack to see details of the routing information. The following shows an example of this information.

```
D TCPIP,TCPCS,NET,ROUTE
EZD0101I NETSTAT CS V1R7 TCPCS
IPV4 DESTINATIONS
DESTINATION         GATEWAY          FLAGS     REFCNT  INTERFACE
193.1.1.94/32       0.0.0.0          H         000000  EZASAMEMVS
193.1.1.94/32       0.0.0.0          UH        000000  EZAXCFC6
193.1.1.94/32       0.0.0.0          UH        000000  EZAXCFC7
193.1.3.94/32       0.0.0.0          UHS       000000  EZAXCFC6
193.1.4.94/32       0.0.0.0          UHS       000000  EZAXCFC7
203.1.1.94/32       0.0.0.0          UH        000000  VIPLCB01015E
213.4.1.1/32        0.0.0.0          UH        000000  LTRLE1A
213.4.2.2/32        0.0.0.0          H         000000  LTRLE2A
213.5.1.1/32        0.0.0.0          UHZ       000001  LTRLE1A
213.6.2.2/32        0.0.0.0          HZ        000001  LTRLE2A
IPV6 DESTINATIONS
DESTIP:   ::1/128
  GW:     ::
  INTF:   LOOPBACK6         REFCNT:  000000
  FLGS:   UH                MTU: 65535
DESTIP:   1EC0::4:1:1/128
  GW:     ::
  INTF:   V6TRLE1A          REFCNT:  000000
  FLGS:   UH                MTU: 14336
DESTIP:   1EC0::4:2:2/128
  GW:     ::
  INTF:   V6TRLE2A          REFCNT:  000000
  FLGS:   H                 MTU: 0
DESTIP:   1EC0::5:1:1/128
  GW:     ::
  INTF:   V6TRLE1A          REFCNT:  000001
  FLGS:   UHZ               MTU: 14336
DESTIP:   1EC0::6:2:2/128
  GW:     ::
  INTF:   V6TRLE2A          REFCNT:  000001
  FLGS:   HZ                MTU: 32000
.
.
32 OF 32 RECORDS DISPLAYED
END OF THE REPORT
```

3. Run the Netstat VCRT/-V DETAIL command on the distributing stack to see the routing information for each connection. The following shows an example of this information.

```
D TCPIP,TCPCS,NET,VCRT,DETAIL
EZD0101I NETSTAT CS V1R7 TCPCS
DYNAMIC VIPA CONNECTION ROUTING TABLE:
Dest:      203.38.1.1..801
  Source:  192.168.2.76..1029
  DestXCF: 193.1.3.94
    PolicyRule:    *NONE*
    PolicyAction:  *NONE*
    Intf:  LTRLE1A
      VipaRoute:  Yes       Gw:  0.0.0.0
Dest:      203.38.1.1..801
  Source:  192.168.2.76..1028
  DestXCF: 193.1.7.94
    PolicyRule:    *NONE*
    PolicyAction:  *NONE*
    Intf:  EZASAMEMVS
      VipaRoute:  No        Gw:  0.0.0.0
Dest:      203.38.1.2..9001
  Source:  192.168.2.76..1031
  DestXCF: 193.1.4.94
    PolicyRule:    *NONE*
    PolicyAction:  *NONE*
    Intf:  LTRLE2A
      VipaRoute:  Yes       Gw:  0.0.0.0
Dest:      203.38.1.2..9001
  Source:  192.168.2.76..1030
  DestXCF: 193.1.6.94
    PolicyRule:    *NONE*
    PolicyAction:  *NONE*
    Intf:  EZASAMEMVS
      VipaRoute:  No        Gw:  0.0.0.0
4 OF 4 RECORDS DISPLAYED
```

Figure 49. Netstat VCRT/-V detail example

Results

See “Routing failures” on page 762 for additional information about routing failures.

Steps for diagnosing Tier 1 z/OS sysplex distribution problems

For an in-depth explanation of distribution to Tier 1 z/OS targets, see the "Sysplex distribution optimizations for multitier z/OS workloads" section in z/OS Communications Server: IP Configuration Guide.

For Tier 1 sysplex distribution

On the Tier 1 distributor, run the Netstat VIPADCFG/-F DETAIL command to see the Tier 1 configuration statement information.

- In the VIPA DEFINE section of the report, verify the new flag field, 1, which indicates that this is a VIPA definition for a Tier 1 DVIPA.
- In the VIPA DISTRIBUTE section, verify the Tier 1 configuration parameters.
  - The flag fields show that the targets are Tier1, and that the distribution method uses Server-specific WLM weights (ServerWLM).
  - The group name is CICSGROUP. The group name is used when a Tier 2 sysplex distributor is being used to correlate Tier 1 targets with their Tier 2 targets. This is explained further in the following Tier 2 distribution.

```
Dynamic VIPA Information:
VIPA Define:
  IpAddr/PrefixLen: 203.1.181.10/24
    Moveable: Immediate  SrvMgr: No  Flg: 1

VIPA Distribute:
  Dest:       203.1.181.10..10000
    DestXCF:   ALL
      SysPt:   No   TimAff: No    Flg: ServerWLM Tier1
      OptLoc:  No
      GrpName: CICSGROUP
      ProcXcost:
        zAAP: 001  zIIP: 001
      ILWeighting: 0
```

Run the Netstat VDPT/-O report to see the status of the tier 1 targets. Verify that weight recommendations are nonzero values and that there is a ready server for each tier 1 target. This information is updated in each polling interval. The ready count is 1 for each tier 1 target, which indicates that each has a server listening on port 10000.

```
Dynamic VIPA Destination Port Table for TCP/IP stacks:
Dest:       203.1.181.10..10000
  DestXCF:   193.1.1.108
  TotalConn: 0000000000  Rdy: 001  WLM: 16  TSR: 100
  Flg: ServerWLM, Tier1
Dest:       203.1.181.10..10000
  DestXCF:   193.1.1.181
  TotalConn: 0000000000  Rdy: 001  WLM: 08  TSR: 100
  Flg: ServerWLM, Tier1
```

Tier 2 sysplex distribution

- A Tier 1 target might distribute connections to Tier 2 targets using a Tier 2 sysplex distributor.
- In the VIPA DEFINE section of the report, the flag field, "2", indicates that this is a VIPA definition for a Tier 2 DVIPA.
- Each Tier 2 distributor will send the weight of each ready server to the Tier 1 distributor; the group name is used to correlate a group of Tier 1 targets with their Tier 2 targets.
- OPTLOCAL has been configured (OptLoc: Yes). As connection requests received are received from the Tier 1 targets, Tier 2 distribution targets on the same stack as the Tier 1 target will be preferred.
- On each Tier 2 distributor run the Netstat VIPADCFG/-F DETAIL report to verify the following:
  - The same group name is used on the Tier 1 and corresponding Tier 2 VIPADISTRIBUTE statements. In the example display from the VIPA distribute section, the group name "CICSGROUP" matches the Tier 1 group name.
  - OPTLOCAL has been configured

```
VIPA Define:
  IpAddr/PrefixLen: 203.2.108.10/24
    Moveable: Immediate  SrvMgr: No  Flg: 2
VIPA Distribute:
  Dest:       203.2.108.10..10002
    DestXCF:   ALL
      SysPt:   No   TimAff: No    Flg: ServerWLM Tier2
      OptLoc:  Yes
      GrpName: CICSGROUP
      ProcXcost:
        zAAP: 001  zIIP: 001
      ILWeighting: 0
```

- On each tier 2 distributor, run the Netstat VDPT/-O report to see the status of each tier 2 target. In the example, there is a nonzero ready count of 1 for the server on destination XCF 193.1.1.108 and the normalized server-specific weight for this target is 16. There is not a ready server on the destination XCF 193.1.1.181. As the tier 1 distributor receives the tier 2 weights, it adds the weight of the corresponding XCF address. The tier 1 normalized weight for the server on destination XCF 193.1.1.108 is twice that of the server weight on destination XCF 193.1.1.181 because the tier 2 distributor did not have a ready server on destination XCF 193.1.1.181.

```
Dynamic VIPA Destination Port Table for TCP/IP stacks:
Dest:        203.2.108.10..10002
  DestXCF:   193.1.1.108
  TotalConn: 0000000000  Rdy: 001  WLM: 16  TSR: 100
  Flg: ServerWLM, Tier2, OptLocal
Dest:        203.2.108.10..10002
  DestXCF:   193.1.1.181
  TotalConn: 0000000000  Rdy: 000  WLM: 00  TSR: 100
  Flg: ServerWLM, Tier2, OptLocal
```

Steps for diagnosing Tier 1 non-z/OS sysplex distribution problems

For an in-depth explanation of distribution to non-z/OS targets, see the Sysplex distribution with DataPower® section in z/OS Communications Server: IP Configuration Guide.

Tip: Tier 1 non-z/OS sysplex distribution is available for both IPv4 and IPv6. The following examples describe how to diagnose IPv4 distribution problems with DataPower, but this function applies to IPv6 as well.

For Tier 1 sysplex distribution

On the tier 1 distributor, run the Netstat VIPADCFG/-F DETAIL report to see the tier 1 configuration statement information.

- In the VIPA DEFINE section of the report, verify the new flag field, 1, which indicates that this is a VIPA definition for a tier 1 DVIPA.
- In the VIPA DISTRIBUTE section, verify the tier 1 configuration parameters.
  - The flag fields show that the targets are Tier1, and that the distribution method is target controlled (TargCtrl). Target controlled distribution means that the tier 1 targets will send weight recommendations to the distributor. The other routing types supported are ROUNDROBIN and WEIGHTEDACTIVE.
  - The routing type used to send packets to the non-z/OS tier 1 targets is GRE. GRE is the only type that is supported.
  - When the RtgType is GRE, a control connection is established from the distributor to each target. In this example, the destination port (CtrlPort) that will be used for the control connection is 1702. The control connection will be used to receive:
    - Weight recommendations from the target (since the distribution method is TargCtrl) - the distributor will load balance across the tier 1 targets using these weights similar to how the distributor uses WLM weight recommendations to load balance across z/OS targets
    - Connection awareness information (since the routing type is GRE) - this information will indicate if the target has a ready server

listening on port 9000, and will also be used to indicate when connections are established and terminated.

– The group name is CICSGROUP. The group name is used when a tier 2 sysplex distributor is being used to correlate tier 1 targets with their tier 2 targets. This is explained further in the following tier 2 distribution section.

```
Dynamic VIPA Information:
VIPA Define:
    IpAddr/PrefixLen: 9.42.130.251/24
      Moveable: Immediate  SrvMgr: No  Flg: 1

VIPA Distribute:
  Dest:        9.42.130.251..9000
    DestXCF:   9.42.105.53
      SysPt:   No   TimAff: No    Flg: TargCtrl Tier1
      OptLoc:  No
      GrpName: CICSGROUP       RtgType: GRE CtrlPort: 01702
  Dest:        9.42.130.251..9000
    DestXCF:   9.42.105.73
      SysPt:   No   TimAff: No    Flg: TargCtrl Tier1
      OptLoc:  No
      GrpName: CICSGROUP       RtgType: GRE CtrlPort: 01702
  Dest:        9.42.130.251..9000
    DestXCF:   9.42.103.215
      SysPt:   No   TimAff: No    Flg: TargCtrl Tier1
      OptLoc:  No
      GrpName: CICSGROUP       RtgType: GRE CtrlPort: 01702
  Dest:        9.42.130.251..9000
    DestXCF:   9.42.103.216
      SysPt:   No   TimAff: No    Flg: TargCtrl Tier1
      OptLoc:  No
      GrpName: CICSGROUP       RtgType: GRE CtrlPort: 01702
```

Run the Netstat ALLCONN/-a report to verify that a control connection is active to each non-z/OS target. The foreign socket destination IP address and destination port correspond to the DestXCF IP address (of the non-z/OS target) and the control port of 1702 from the VIPA distribute section. In this example report, there are four TCP connections in established state to each of the configured Tier 1 targets.

```
USER1@VIC018:/:> netstat -a
MVS TCP/IP NETSTAT CS V1R11       TCPIP Name: TCPCS           17:23:36
User Id  Conn     State
-------  ----     -----
TCPCS    00000014 Establsh
  Local Socket:   9.42.104.4..1025
  Foreign Socket: 9.42.103.216..1702
TCPCS    00000018 Establsh
  Local Socket:   9.42.104.4..1027
  Foreign Socket: 9.42.105.73..1702
TCPCS    00000019 Establsh
  Local Socket:   9.42.104.4..1028
  Foreign Socket: 9.42.105.53..1702
TCPCS    00000017 Establsh
  Local Socket:   9.42.104.4..1026
  Foreign Socket: 9.42.103.215..1702
```

Run the Netstat VDPT/-O report to see the current status of the Tier 1 targets. Verify the weight recommendations are non-zero and that there is a ready server for each Tier 1 target. This information is updated each polling interval.

- The T1Wt field is the weight recommendation reported by each target. In the example, all are non-zero.

- The ready count is 1 for each Tier 1 target indicating that each has a server Listening on port 9000.
- The CWt field is a CPC weight. This field should be non-zero if Tier 2 distributors are configured. If there are no corresponding Tier 2 distribution statements, then this field will be zero. This is explained further in the following Tier 2 distribution section.

```
Dynamic VIPA Destination Port Table for non-z/OS targets:
Dest:        9.42.130.251..9000
  Target Addr: 9.42.103.215
  TotalConn: 0000000000  Rdy: 001  Wt: 03 CWt: 032
  Flg: TargCtrl
    T1Wt: 420
    ActConn:   0000000000
Dest:        9.42.130.251..9000
  Target Addr: 9.42.103.216
  TotalConn: 0000000000  Rdy: 001  Wt: 05 CWt: 032
  Flg: TargCtrl
    T1Wt: 672
    ActConn:   0000000000
Dest:        9.42.130.251..9000
  Target Addr: 9.42.105.53
  TotalConn: 0000000000  Rdy: 001  Wt: 03 CWt: 064
  Flg: TargCtrl
    T1Wt: 420
    ActConn:   0000000000
Dest:        9.42.130.251..9000
  Target Addr: 9.42.105.73
  TotalConn: 0000000000  Rdy: 001  Wt: 03 CWt: 064
  Flg: TargCtrl
    T1Wt: 415
    ActConn:   0000000000
```

Tier 2 sysplex distribution

- A Tier 1 target may distribute connections to Tier 2 targets on the same CPC using a Tier 2 sysplex distributor configured with CPCSCOPE. The Tier 2 sysplex distributor will only load balance to z/OS targets that are on the same CPC.
- In the VIPA DEFINE section of the report, the new flag field, 2, indicates that this is a VIPA definition for a Tier 2 DVIPA. The flag, "C", indicates that the Tier 2 DVIPA is configured with CPCSCOPE.
- Each Tier 2 distributor will send a combined CPC weight of all ready servers to the Tier 1 distributor; the group name is used to correlate a group of Tier 1 targets with their Tier 2 targets.
- On each Tier 2 distributor run the Netstat VIPADCFG/-F DETAIL report to verify that the same group name is used on the Tier 1 and corresponding Tier 2 VIPADISTRIBUTE statements. In the example report from the VIPA distribute section, the group name "CICSGROUP" matches the Tier 1 group name.
- A combined weight will be used to update a Tier 1 target with a matching group name if this Tier 2 distributor (and therefore its Tier 2 targets) are on the CPC used by that Tier 1 target. A VIPADEFINE CPCSCOPE definition is required to determine the CPC used by a Tier 1 target. If the IP address of the target is in the same subnet of this DVIPA, then its targets are on that CPC.
  - In the example report from the VIPA define section, DVIPA 9.42.103.17 is a CPCSCOPE DVIPA because the flag field is "C".

- The subnet for this DVIPA is 9.42.103.xx/24. So all Tier 1 targets (from the Tier 1 Netstat VDPT/-O report above) in this subnet are using the same CPC as that stack.
- The Tier 1 targets in this subnet are 9.42.103.215 and 9.42.103.216. So as combined weights are received from a Tier 2 distributor, with the same CPC ID as the stack in which the CPCSCOPE definition was defined, the Tier 1 combined weight fields (CWt) for those targets (in the CPCSCOPE subnet) are updated with those weights.

- In this example, the CPCSCOPE DVIPA is configured on the same stack as the Tier 2 VIPADISTRIBUTE statement. So this Tier 2 distributor will be providing the combined weights for the non-z/OS targets in the CPCSCOPE DVIPA subnet, 9.42.103.215 and 9.42.103.216. Ideally the CPCSCOPE DVIPA would be defined on the same stack as the Tier 2 Distributor, and VIPABACKUP statements for these DVIPAs would be on the same stacks with the same ranks. However the CPCSCOPE DVIPA could be defined on any stack as long as it is in the same CPC as the Tier 2 distributor.

```
VIPA Define:
  IpAddr/PrefixLen: 9.42.103.17/24
    Moveable: Immediate  SrvMgr: No  Flg: C
  IpAddr/PrefixLen: 203.2.108.3/24
    Moveable: Immediate  SrvMgr: No  Flg: 2C
VIPA Distribute:
  Dest:        203.2.181.3..11000
    DestXCF:   ALL
      SysPt:   No   TimAff: No    Flg: ServerWLM Tier2
      OptLoc:  No
      GrpName: CICSGROUP
```

- On each Tier 2 distributor, run the Netstat VDPT/-O DETAIL report to see the status of each Tier 2 target. In the example, there is a non-zero ready count of 1 for each server, and there is a non-zero normalized server-specific weight for each target.
- The combined weight in the Tier 1 VDPT (CWt) displays for the Tier 1 targets on that CPC, 9.42.103.214 and 9.42.103.216 is the total weight of the Tier 2 ready servers, 32.

```
Dynamic VIPA Destination Port Table for TCP/IP stacks:
Dest:        203.2.108.3..11000
  DestXCF:   193.1.1.108
  TotalConn: 0000000000  Rdy: 001  WLM: 16  TSR: 100
  Flg: ServerWLM, Tier2
    TCSR: 100  CER: 100 SEF: 100
    Weight: 64
      Raw          CP: 64 zAAP: 00 zIIP: 00
      Proportional CP: 63 zAAP: 00 zIIP: 00
    Abnorm: 0000       Health: 100
    ActConn:   0000000000
    QosPlcAct: *DEFAULT*
      W/Q: 16
Dest:        203.2.108.3..11000
  DestXCF:   193.1.1.181
  TotalConn: 0000000000  Rdy: 001  WLM: 16  TSR: 100
  Flg: ServerWLM, Tier2
    TCSR: 100  CER: 100 SEF: 100
    Weight: 64
      Raw          CP: 64 zAAP: 00 zIIP: 00
      Proportional CP: 63 zAAP: 00 zIIP: 00
    Abnorm: 0000       Health: 100
    ActConn:   0000000000
    QosPlcAct: *DEFAULT*
      W/Q: 16
```

Chapter 12. Diagnosing access control problems

This topic describes selected procedures for TCP/IP Services component trace, packet trace, and Socket API trace.

This topic contains the following sections:

- "Overview of access control support" on page 418
- "Diagnosing multilevel security consistency check messages (EZD1215-EZD1234)" on page 419

Tip: The SAF interface allows security servers to return the following responses to access control questions:

Allow User is permitted to resource with requested level of access.

Deny User is not permitted to resource with requested level of access.

No decision
Class is not active or covering profile is not defined.

For many resources, TCP/IP allows access when a No decision is returned. RACF supports the No decision response. Some security server products do not support the No decision response. They always return Deny when a resource has no profile. If you are using one of these other security servers, you must define profiles for these resources to allow any user to use them.
TCP/IP creates resource names in the SERVAUTH class to represent the services it protects.

These resource names are comprised of the following tokens:

- The first token is always EZA or EZB.
- The second token represents the type of services.
- The third token is the eight-character MVS system image name.
- The fourth token is often the TCP/IP job name.

Additional tokens can be defined for more granularity on certain types of services. For more information about services that TCP/IP protects and the resource names used, see the security topic in z/OS Communications Server: IP Configuration Guide.

You define RACF profiles in the SERVAUTH class to control access permissions to these resource names. A discrete profile has the same name as a resource and covers only that resource. A generic profile uses wildcard symbols to cover many resource names. The SERVAUTH class is a general resource class, so you use the RACF RDEFINE, RLIST, RALTER, RDELETE and PERMIT commands to manage these profiles. For more information, see z/OS Security Server RACF Security Administrator's Guide and z/OS Security Server RACF Command Language Reference.

Except for a few documented cases, TCP/IP checks for READ access to resources. Users might be given access to a resource in several ways. A RACF profile defines universal access (UACC) that provides the default level of access for all users not explicitly named. Individual users and user groups might be given a different level

© Copyright IBM Corp. 2000, 2015

of access, higher or lower, with the PERMIT command. Use the WHEN clause to define conditions that must be met before the specified access is granted.

Tip: The RACF WHEN(PROGRAM(...)) clause has restrictions on profiles in the SERVAUTH class. It can be ignored on some resource checks and should only be used for resource names that explicitly document support.

RACF can be configured to write audit messages to the console. The default for profiles in the SERVAUTH class is to write a message when access is denied. These messages indicate the user, resource name, profile name and access level requested. When you first put an access control policy in place, you might want to configure the profile to produce audit messages on successes as well. You might also want to configure the profile with the WARNING parameter. This causes RACF to write the audit failure messages and then return allow to the resource manager. This allows you to test the effectiveness of a proposed policy without impacting usage.

Tips:

- Some policy changes do not take effect until the next time a user logs on or starts a job. After changing the policy, the user might need to log off or a job might need to be canceled and restarted.
- TCP/IP caches results when it checks access to NETACCESS resources. This cache is purged when a NETACCESS statement is found in a file used with the VARY TCPIP,,OBEYFILE command. It is also purged when an ENF signal is received from RACF indicating that the SERVAUTH class or SECLABEL class has been refreshed. If your security server does not produce this ENF signal then, after making policy changes, you must issue the VARY TCPIP,,OBEYFILE command with a file containing the NETACCESS statement to cause TCP/IP to purge cached responses.

Several of the TCP/IP services that provide access control check socket calls made through several different interfaces. When access to a resource is denied, the errno returned is EACCES. The errno2 field provides additional information about the failure. Programs that provide diagnostic logs should include the errno2 field. For information about the contents of the value returned, see z/OS UNIX System Services Programming: Assembler Callable Services Reference.

Tip: Many C programs use the perror() or strerror() library service to display errors encountered. There is an environment variable _EDC_ADD_ERRNO2, which when set to 1, appends the current errno2 value to the end of the perror() string as shown below:

```
EDC5121I Invalid argument. (errno2=0x0C0F8402)
```

TCP/IP access control failures are recorded in the event trace (SYSTCPIP) for TCP/IP stacks with the ACCESS option.

Overview of access control support

Communications Server is a resource manager that provides access control support over many of its services.

This can be a powerful tool to prevent unwanted usage of communications services. At times, it might also prevent intended usage. TCP/IP uses Security Access Facility (SAF) interfaces to ask your installed security server access control questions.

Note: The examples and terminology in this topic assume you are using RACF. However, you can use any SAF-conforming security server.

Diagnosing multilevel security consistency check messages (EZD1215-EZD1234)

Secure communication in a multilevel secure environment requires configuration of several statements in the TCPIP.PROFILE and security server resource profiles in the SERVAUTH, SECLABEL and STARTED classes. Inconsistencies in this configuration can allow unintended communication or prevent intended communication. When the RACF MLACTIVE option is set, TCP/IP checks the TCPIP.PROFILE and security server resource profiles for consistency. Consistency checking occurs at TCP/IP initialization, when a VARY TCPIP,,OBEYFILE command is processed and when RACF sends an ENF signal specifying that a RACLIST REFRESH was done on the SERVAUTH or SECLABEL class.

TCP/IP writes an informational message to the job log for each inconsistency detected. If inconsistencies are found, a final message, EZD1217I, summarizing the number of problems found is written to the system console. You should check the job log for messages in the range EZD1219I-EZD1234I whenever message EZD1217I appears on the system console. You should correct your configuration as indicated by the job log messages until TCP/IP no longer detects any errors.

TCP/IP's default behavior is to continue running when inconsistent security configurations are detected. If you plan to run in a multilevel-secure environment, it is recommended that you specify GLOBALCONFIG MLSCHKTERMINATE in your TCPIP.PROFILE when running production workloads and GLOBALCONFIG NOMLSCHKTERMINATE while you are making planned changes to your security environment.

Steps for verifying the configuration

Verify that the TCP/IP stack is configured as intended.

Before you begin

See z/OS Communications Server: IP Configuration Guide for information about networking in a multilevel-secure environment.

Procedure

Perform the following steps:

1. TCP/IP stack is running under the intended user ID. If the stack is a submitted job, check the USER= parameter on the job card. If the stack is a started procedure, check the STDATA segment of the profile in the STARTED class.
2. TCP/IP stack is running with the intended security label. If the stack is a submitted job, check the SECLABEL= parameter on the job card. If the stack is a started procedure or SECLABEL= was not specified on the job card, check the default security label in the USER profile. Verify that the user ID is permitted to the SECLABEL profile. If running with the RACF SECLBYSYSTEM option, verify that the security label is active on this system image.
3. TCP/IP stack recognizes the multilevel-secure environment. The TCPIP.PROFILE must contain a valid NETACCESS statement with the following:
   - INBound

- OUTBound
- At least one valid security zone definition

4. TCP/IP stack has the intended IP addresses defined. Verify the IP addresses on DEVICE and INTERFACE statements in the TCPIP.PROFILE. Verify the IP addresses on VIPADEFINE, VIPABACKUP, VIPARANGE and VIPADISTRIBUTE statements in the TCPIP.PROFILE. Verify that IP addresses are manually configured for IPv6 interfaces. Verify that the INTFID keyword is specified on all IPv6 interfaces. Verify that the IPADDR keyword is specified on all IPv6 interfaces that support autoconfiguration.
5. TCP/IP stack has IP addresses mapped into the intended network security zones. Verify that the base IP address, mask and zone name are correct on each line in NETACCESS statement in the TCPIP.PROFILE. Verify that these addresses are in security zones:
   - INADDR_ANY (IPv4 0.0.0.0/32, IPv6 ::/128)
   - LOOPBACK (IPv4 127.0.0.1/8, IPv6 ::1/128)
   - Any required Multicast (IPv4 224.0.0.0/4, IPv6 FF00::/8)

   Tips:
   - The D TCPIP,,N,ACC,NETW console command displays the current NETACCESS statement configuration. The SERVAUTH profile name covering the security zone resource name and the security label defined on that profile are also shown.
   - The security zone that a IP address is currently configured into is displayed by the D TCPIP,,N,ACC,NETW,ipaddress console command.
6. SERVAUTH resources are covered by the intended profile. The RLIST SERVAUTH resource_name AUTHUSER RACF command displays the discrete or generic profile that most closely matches the specified resource name. It also displays the universal access, the security label, the access list and the conditional access list for that profile.

Chapter 13. Diagnosing line print requester and daemon (LPR and LPD) problems

Line print requester (LPR) and line printer daemon (LPD) compose a functional unit in which the LPR client sends data files to a printer controlled by an LPD server. These files can be in ASCII form or extended binary-coded decimal interchange code (EBCDIC) form.

In most environments, customers have different types of LPR clients and LPD servers, running on platforms, such as MVS, OS/2, AIX®, and UNIX. However, all print client and servers must follow the standards contained in RFC1179. Some clients and servers provide more than what is required by the RFC, while some clients and servers are restricted or limited, which can cause errors or require more configuration to work.

On platforms, such as MVS, UNIX, and AIX, you can start the LPR client program with command prompts, through batch (in MVS), or through shell scripts (in UNIX/AIX®). The MVS LPD server allocates temporary data sets to process incoming print requests from various clients. These data sets use the TCP/IP high-level qualifiers (HLQs) or the prefix defined in the LPD server cataloged procedure.

The MVS LPD server can also act as a client when a remote print server is defined in the LPD configuration file as a *service*. In this case, when the LPD server receives an incoming print job, it opens a new connection through a client port, and sends the data to the remote print server. When a remote print server is used, LPD specifications, such as line size and page size, do not apply. Instead, the specifications of the remote server apply.

For information about configuring your LPD server, see z/OS Communications Server: IP Configuration Reference. For information about using the client-related LPR, LPQ, and LPRM commands, see z/OS Communications Server: IP User's Guide and Commands.

Problems with the print function are easy to diagnosis if the problem is within the LPR client or the LPD server. More difficult problems can be encountered in the TCP/IP layer or in sockets. In addition, incorrectly built or defined translation tables can produce unpredictable results, such as abends, misprinted characters, and hang conditions (caused by delayed acknowledgments).

Diagnosing LPR client and LPD server problems

Problems with LPR and LPD generally fall into one of the following categories:

- Abends
- Timeouts, hangs, and waits
- Incorrect output

These categories are described in the following subtopics.

© Copyright IBM Corp. 2000, 2015

Abends

When an abend occurs during LPD processing, messages and other error-related information are sent to the MVS system console. If this information is insufficient to solve the problem, use the information provided in a dump. To produce a dump, code a SYSMDUMP DD or SYSABEND DD statement in the LPD cataloged procedure. If you do not do the coding before the abend occurs, code the statement after the abend, re-create the abend or wait for it to occur again. For information about analyzing dumps produced during LPD processing, see z/OS Problem Management.

It can also be helpful to obtain and analyze information from the following sources:

- LPD trace in the SYSPRINT data set
- Output of LPD started task
- System log (syslog)

Steps for diagnosing timeouts, hangs, and waits

Timeouts, hangs, and waits occur when the LPD server does not respond to client requests for a data packet, an acknowledgment that a data packet was received, or a reply to a command. Similarly, the LPD server can time out a connection if the LPR client does not respond.

Before you begin

Determine whether one or more of the following problems caused a timeout, hang, or wait:

- Incorrect host name or IP address specified on the LPR command
- Malfunctioning remote server or remote host
- Problems with the network (for example, network congestion), bridge, gateway, or router in the routing path
- Problems with the device or channel attached to the host
- Corrupted TCP/IP address space
- Incorrectly built or defined translation tables
- Malfunctioning LPR client

Procedure

Perform one or more of the following steps to diagnose timeouts, hangs, and waits.

1. Check to see whether the target LPD print server is running, has enough paper, and is not jammed.
2. Check the LPR and LPD traces for possible error messages, or for the last activity performed by LPR or LPD (for example, waiting for a connection, port availability, or an acknowledgment). Be aware that when sending a print request to a remote printer through the LPD server, the LPR client can show a successful data transfer even though there might be a problem connecting to the remote printer.
3. Check the IP address or host name used with the LPR command.
4. Check the LPR, LPD, and packet traces. If the packet trace shows a problem during binding or connecting, then check the socket trace.
5. Verify that the translation tables are built correctly. Test them using the *hlq*.STANDARD.TCPXLBIN table supplied with TCP/IP.

Results

Be aware that waits can occur because some LPD servers do not send acknowledgments until data is printed. In this situation, the LPR client does not show successful data transfer until it receives the acknowledgment.

Incorrect output

LPR problems with incorrect output fall into one of the following categories:

- Garbled data sent from the LPR client or received by the LPD server
- Truncated or missing print data
- LPR works with some options, but not others

These categories are described in the following topics.

Steps for diagnosing garbled data

Determine why data is not properly translated or transferred.

Procedure

Perform the following steps:

1. Determine whether the binary option or the default EBCDIC was used when the data file was printed. If the binary option was used, the LPR client did not translate the data. If EBCDIC was used, check for erroneous control characters or conflicting combinations of options.
2. Check to see whether other files print correctly from the same client and to the same server. Check to see whether the problem file prints correctly to other servers.
3. Verify that the translate tables for the sender and receiver are reciprocals of each other. Determine which characters are consistently garbled and examine those entries in the tables. To determine the name of the translation table used by the LPR client, check the LPR messages issued at startup.
4. Check the IP packet trace to determine exactly what data was sent from the client and acknowledged by the LPD server.
5. If data shown in the IP packets from the LPR client to the server is correct, there might be an error on the server or printer. Check the server traces and setup on the printer or LPD server. Some servers require certain printer names or options to be specified on the LPR (lp from omvs) commands.

Steps for diagnosing truncated or missing print data

Determine why print data is truncated or missing.

Procedure

Perform one or more of the following steps:

1. Check to see whether the value for the record length is valid. The value is specified using the WIDTH option and variable on the LPR command.
2. If MVS displays truncated records, check the value of the LINESIZE option on the SERVICE statement in the LPD configuration file.
3. If you use the FILTER L or FILTER R options on the LPR command, check to see whether the control characters on the first column of the source file are valid. LPR issues a message indicating whether a record of data has been ignored.

4. Using a packet trace and the file size listed in the LPR trace control record, verify that the correct number of bytes were sent by the LPR client and received by the LPD server.
5. Check the LPD trace for error messages. Verify that the `Job xxx Received` and `Job xxx Finish Printing` messages were received.
6. If sending a print request to a remote printer through the LPD server, check the LPD trace to determine whether all data were sent successfully to the remote printer. If not or if data are incorrect, check the printer for errors or restrictions on the type of data it supports (for example, postscript only, text only).
7. Check for partial temporary data sets and either rename them or delete them. The LPD server creates temporary data sets when connections are broken, and the server does not completely process a print job. (Depending on the LPR client, the server can requeue the job for printing at a later time.) When the connection is restored, the daemon checks for temporary data sets and processes them. After processing, they are erased.

   The temporary data sets are stored on a volume with a data set prefix you define in the LPD cataloged procedure. Following are samples of these data sets:

```
TCPUSR4.PRT1.QUEUE                                      WRKLB2
TCPUSR4.RALVM12.CFnnn                    BROWSED        WRKLB2
TCPUSR4.RALVM12.DFAnnnLU                 BROWSED        TCPWRK
TCPUSR4.RALVM12.JFnnn                                   WRKLB2
```

```
The QUEUE... represents, in this sample PRT1's print
             queue file.  It will contain the name of
             the JOB files that have not been completely
             processed yet.
The CF... represents the CONTROL FILE.
          Contains the control data/commands sent to LPD.
The DF... represents the DATA FILE.
          The actual data sent to be printed.
The JF... represents the JOB FILE.
          Contains names of the above files that have
          not been processed yet.
where nnn is the three digit job number.
```

   Occasionally, depending on the precipitating incident and the time the connection was broken, the LPD server creates only a portion of one or more data sets. When partial temporary data sets are created, the server issues allocation or failure-to-erase messages. If you receive any of these messages, search for the partial data sets and either rename or delete them. After doing this, you might need to reissue the print request or requests.

   The LPD trace and the system log at the time a connection is broken show the status of all print jobs (and the status of some data sets) and identify the owners of the print requests.

Steps for diagnosing LPR working with some options only

Determine why the LPR command works with some options but not with others.

Procedure

Perform one or more of the following steps:

1. If some print requests do not work with certain LPR options, check the LPR trace for error messages.
2. If the LPR command from batch fails, but works under TSO, check for possible errors in the batch-job output and for error messages in the LPR trace.

Results

For information about the LPR command, see z/OS Communications Server: IP User's Guide and Commands.

LPR client traces

This topic provides information about activating LPR client traces. It also provides samples of trace output with explanations of selected messages.

Step for activating LPR client traces

You can activate LPR client traces by specifying the TRACE option in addition to the usual processing parameters on the LPR command.

For example, enter the following command to start the LPR client with trace on:

```
LPR filename (Printer p1 Host h1 TRACE
```

Step for creating client trace output

LPR trace output is sent to SYSOUT and can be displayed on the LPR client console. Figure 50 on page 426 is a sample of an LPR trace created by way of TSO with the following command:

```
LPR soto.files(lpconfig) (p prt1 h 9.67.113.60 TRACE
```

```
EZB0915I Begin "LPR" to printer "prt1" at host "9.67.113.60"
 1
EZB1057I Loaded translation table from "TCP31S.STANDARD.TC PXLBIN".
EZB0920I Requesting TCP/IP service at 96155 18:52:53
EZB0921I Granted TCP/IP service at 96155 18:52:53
EZB0922I Resolving 9.67.113.60 at 96155 18:52:53
EZB0924I Host 9.67.113.60 name resolved to 9.67.113.60 at 96155 18:52:53
EZB0925I TCP/IP turned on.
EZB0926I Host "MVSA" Domain "TCP.RALEIGH.IBM.COM" TCPIP Service Machine TCP31S
EZB0927I Trying to open with local port 721 to foreign host address 9.67.113.60
 2
EZB0928I Connection open from local port 721 to foreign host address 9.67.113.60
EZB0961I Control file name is cfA827MVSA
EZB0962I Data file name is dfA827MVSA Port Number=721. Remote IP Addr=9.7.113.60
 3
EZB0916I Sending command 2 argument: prt1  Port Number=721. Remote IP Addr=9.67.113.60
EZB0917I Command successfully sent  Port Number=721. Remote IP Addr=9.67.113.60
EZB1012I Receiving ACK Port Number=721. Remote IP Addr=9.6 7.113.60
EZB1013I ReceiveACK: TRUE for byte value 00 Port Number=721. Remote IP Addr=9.67.113.60
EZB0997I Byte size check starts at 96155 18:52:54
EZB0998I Byte size check ends at 96155 18:52:54
EZB0999I Send command starts at 96155 18:52:54 Port Number=721.   Remote IP Addr=9.67.113.60
 4
EZB0916I Sending command 3 argument:7434 dfA827MVSA Port Number=721. Remote IP Addr=9.67.113.60
EZB0917I Command successfully sent Port Number=721. Remote IP Addr=9.67.113.60
 5
EZB1012I Receiving ACK  Port Number=721. Remote IP Addr=9.67.113.60
 5
EZB1013I ReceiveACK: TRUE for byte value 00 Port Number=721. Remote IP Addr=9.67.113.60
EZB1000I Send command ends at 96155 18:52:55 Port Number=721. Remote IP Addr=9.67.113.60
 6
EZB1001I Send data starts at 96155 18:52:55  Port Number=721. Remote IP Addr=9.67.113.60
 6
EZB1002I Send data ends at 96155 18:52:56 Port Number=721. Remote IP Addr=9.67.113.60
EZB1003I Send ACK starts at 96155 18:52:56 Port Number=721. Remote IP Addr=9.67.113.60
EZB1014I Sending ACK Port Number=721. Remote IP Addr=9.67. 113.60
 7
EZB1015I ACK successfully sent Port Number=721. Remote IP Addr=9.67.113.60
EZB1004I Send ACK ends at 96155 18:52:56 Port Number=721. Remote IP Addr=9.67.113.60
EZB1012I Receiving ACK Port Number=721. Remote IP Addr=9.6 7.113.60
 8
EZB1013I ReceiveACK: TRUE for byte value 00 Port Number=721. Remote IP Addr=9.67.113.60
 9
EZB1009I Data file sent. Port Number=721. Remote IP Addr=9.67.113.60
EZB1011I Queuing control line "HMVSA.TCP.RALEIGH.IBM.COM"
EZB1011I Queuing control line "PTCPUSR4"
EZB1011I Queuing control line "JTCPUSR4.SOTO.FILES(LPCONFIG)"
EZB1011I Queuing control line "CMVSA.TCP.RALEIGH.IBM.COM"
EZB1011I Queuing control line "LTCPUSR4"
```

Figure 50. Example of LPR trace output (Part 1 of 2)

```
 10
EZB1011I Queuing control line "fdfA827MVSA"
EZB1011I Queuing control line "UdfA827MVSA"
EZB1011I Queuing control line "NTCPUSR4.SOTO.FILES/LPCONFIG"
 11
EZB0916I Sending command 2 argument: 153 cfA827MVSA Port Number=721.   Remote IP Addr=9.67.113.60
EZB0917I Command successfully sent Port Number=721. Remote IP Addr =9.67.113.60
EZB1012I Receiving ACK  Port Number=721. Remote IP Addr=9.6 7.113.60
 12
EZB1013I ReceiveACK: TRUE for byte value 00 Port Number=721. Remote IP Addr=9.67.113.60
 13
EZB1017I Control data sent Port Number=721. Remote IP Addr=9.67.113.60
EZB1014I Sending ACK Port Number=721. Remote IP Addr=9.67. 113.60
EZB1015I ACK successfully sent Port Number=721.   Remote IP Addr=9.67.113.60
EZB1012I Receiving ACK Port Number=721. Remote IP Addr=9.6 7.113.60
 14
EZB1013I ReceiveACK: TRUE for byte value 00  Port Number=721. Remote IP Addr=9.67.113.60
 15
EZB1018I Control file sent  Port Number=721. Remote IP Addr=9.67.113.60
```

Figure 51. Example of LPR trace output (Part 2 of 2)

Following are short descriptions of the numbered items in the trace:

1 Indicates the translation table used by the LPR client. In this print request, no translation tables were defined by the person submitting the request.

2 Indicates LPR port used to connect to the LPD server with the IP address 9.67.113.60. The LPR port range is from 721 through 731.

3 Indicates the LPR command sent to the LPD server identifying the name of the print queue where the output was sent. See RFC1179 for details on commands and subcommands issued between LPR and LPD.

4 Indicates the command that provided the LPD print server with the byte size (7434) and name of the data file (dfA827MVSA) that was sent.

- The character string dfA indicates that this was a data file.
- The number 827 was the three-digit job number that was randomly generated by the LPR client or specified in the LPR command using the **JNUM** option.
- MVSA was the name of the host from which the print request came.

5 Indicates the client is waiting for the LPD server to acknowledge the sending command in item **4** . The message on the following line (`TRUE (00)`) indicates that the client received an acknowledgment. A `FALSE` message or any value other than zero terminates the LPR print request.

6 Indicates that the LPR client started and then stopped sending the data file.

7 Indicates that the LPR client notified the LPD server, by way of an acknowledgment, that the complete file was sent. The LPR client waits for the server to acknowledge receipt of the entire data file.

8 Indicates that the client received an acknowledgment from the server that the entire data field was received.

9 Confirms that the data file was sent to the LPD server.

10 Specifies one of the several control records sent by the LPR client. (The records are described in detail in RFC1179.) This control record is mandatory and represents the name of the data file created by the LPD server. The name is preceded by the filter specified on the LPR command. The letter **f** denotes the default filter.

11 Specifies the byte size (153) and the name of the control file (cfA827MVSA) that was sent.

12 Indicates that the LPD server received the command and expected the control file to be sent.

13 Indicates that the LPR client sent the control file and an acknowledgment that it finished sending the entire file. The last line in the block indicates that the client was waiting for an acknowledgment from the server.

14 `TRUE (00)` indicates that the client received an acknowledgment from the LPD server that the control file was received.

15 Confirms that the control file was sent to the LPD server. The job was then terminated.

Figure 52 is a sample LPR trace showing a print request in which the FILTER X option was specified on the LPR command. Since the LPD server does not support this type of filter, it rejects the print request. (For an example of an LPD trace that shows that this job was rejected, see Figure 57 on page 437.) The LPR trace does not show an error because it can send a print request to non-IBM LPDs that support other filters (for example, FILTER X). For detailed information about filters, see RFC1179 and to z/OS Communications Server: IP Configuration Reference.

The trace was produced using the following command issued through TSO by user ID TCPUSR4:

```
LPR test (p TIANNA h 9.67.113.60 filter x TRACE
```

```
1
EZB0915I Begin "LPR" to printer "TIANNA" at host "9.67.113.6  0"
EZB1057I Loaded translation table from "TCP31S.STANDARD.TCPXLBIN".
EZB0920I Requesting TCP/IP service at 96155 19:22:15
EZB0921I Granted TCP/IP service at 96155 19:22:15
EZB0922I Resolving 9.67.113.60 at 96155 19:22:15
EZB0924I Host 9.67.113.60 name resolved to 9.67.113.60 at 96155 19:22:15
EZB0925I TCP/IP turned on.
EZB0926I Host "MVSA" Domain "TCP.RALEIGH.IBM.COM" TCPIP Service Machine TCP31S
EZB0927I Trying to open with local port 721 to foreign host address 9.67.113.60
EZB0928I Connection open from local port 721 to foreign host address 9.67.113.60
:
2
EZB1009I Data file sent.  Port Number = 721.   Remote IP Addr  = 9.67.113.60
3
EZB1011I Queuing control line "HMVSA.TCP.RALEIGH.IBM.COM"
EZB1011I Queuing control line "PTCPUSR4"
EZB1011I Queuing control line "JTCPUSR4.TEST"
EZB1011I Queuing control line "CMVSA.TCP.RALEIGH.IBM.COM"
EZB1011I Queuing control line "LTCPUSR4"
4
EZB1011I Queuing control line "xdfA947MVSA"
EZB1011I Queuing control line "UdfA947MVSA"
EZB1011I Queuing control line "NTCPUSR4.TEST"
EZB0916I Sending command 2 argument: 122 cfA947MVSA  Port Number = 721. Remote IP Addr = 9.67.113.60
EZB0917I Command successfully sent  Port Number = 721. Remote IP Addr = 9.67.113.60
:
5
EZB1018I Control file sent  Port Number = 721. Remote IP Ad dr = 9.67.113.60
```

Figure 52. Example of LPR trace with filter x option

Following are short descriptions of the numbered items in the trace:

1 Indicates that the print request was issued to a printer named TIANNA at IP address 9.67.113.60.

2 Indicates that the data file was sent. The error was not recognized until the LPD server tried to process the print job. (See Figure 57 on page 437.)

3 Indicates control commands sent to the LPD server. For details about these commands, see RFC1179.

4 Represents the name of the data file. The character string `xdf` indicates that the x filter was used.

5 Indicates that the control file was sent to the LPD server. The job was then terminated.

Figure 53 is a sample showing a print request using the following command `lpr test (p njeSOTO host MVSA` without the TRACE option. The output shows an error because the printer name was not entered entirely in capital letters.

```
1
EZB1006E Host MVSA did not accept printer name njeSOTO.
             Port Number = 721  Remote IP Addr = 9.67.113.60
2
EZB1049E Send printer command did not receive ACK.  ACK message = .
             Port = 721.   Remote IP Addr = 9.67.113.60
```

Figure 53. Example of LPR output with unknown printer

Following are short descriptions of the numbered items in the trace.

1 Indicates that a SERVICE statement for a printer named njeSOTO did not exist in the LPD server configuration file.

2 Indicates that the LPD server did not send a positive response to the LPR client. The job was then terminated.

Figure 54 on page 430 is a sample LPR trace output produced with the following command the JNUM option and variable, along with the LANDSCAPE and TRACE options:

```
lpr test (p TIANNA host 9.67.113.60 JNUM 111 LANDSCAPE TRACE
```

The trace output shows the scanning that occurred to identify the first available port.

```
 1 EZB0988I PostScript program is 635 bytes
EZB0915I Begin "LPR" to printer "TIANNA" at host "9.67.113.60"
EZB1057I Loaded translation table from "TCP31S.STANDARD.TCPXLBIN".
EZB0920I Requesting TCP/IP service at 96155 19:35:12
EZB0921I Granted TCP/IP service at 96155 19:35:12
EZB0922I Resolving 9.67.113.60 at 96155 19:35:12
EZB0924I Host 9.67.113.60 name resolved to 9.67.113.60 at 96155 19:35:12
EZB0925I TCP/IP turned on.
EZB0926I Host "MVSA" Domain "TCP.RALEIGH.IBM.COM" TCPIP Service Machine TCP31S
 2 EZB0927I Trying to open with local port 721 to foreign host a ddress 9.67.113.60
EZB0927I Trying to open with local port 722 to foreign host address 9.67.113.60
EZB0927I Trying to open with local port 723 to foreign host address 9.67.113.60
EZB0927I Trying to open with local port 724 to foreign host address 9.67.113.60
 3 EZB0928I Connection open from local port 724 to foreign host  address 9.67.113.60
 4 EZB0961I Control file name is cfA111MVSA
EZB0962I Data file name is dfA111MVSA  Port Number = 724.   Remote I P Addr = 9.67.113.60
EZB0916I Sending command 2 argument: TIANNA  Port Number = 724.   Remote IP Addr = 9.67.113.60
:
EZB1009I Data file sent.  Port Number = 724.   Remote IP Addr = 9.67 .113.60
EZB1011I Queuing control line "HMVSA.TCP.RALEIGH.IBM.COM"
EZB1011I Queuing control line "PTCPUSR4"
EZB1011I Queuing control line "JTCPUSR4>TEST"
EZB1011I Queuing control line "CMVSA.TCP.RALEIGH.IBM.COM"
EZB1011I Queuing control line "LTCPUSR4"
 5 EZB1011I Queuing control line "fdfA111MVSA"
EZB1011I Queuing control line "UdfA111MVSA"
EZB1011I Queuing control line "NTCPUSR4.TEST"
EZB0916I Sending command 2 argument: 122 cfA111MVSA  Port Number = 7 24.
```

Figure 54. Example of LPR trace with JNUM, LANDSCAPE, and TRACE options

Following are short descriptions of the numbered items in the trace:

1 Indicates that the LPR client inserted a landscape header, written in postscript, at the beginning of the data file.

2 Indicates that the LPR client was attempting to use the first available client port. The port range for the LPR client is 721 through 731. If no ports are available, an error message is displayed.

3 Indicates that a connection was opened using port 724.

4 Indicates that the value specified for JNUM (111) was used to build the control and data file names.

5 Indicates the name of the file containing the three-digit job number that was used with the file name sent to the print server.

Following is a clipping of the header that was inserted into the data file. For more information about header files, see z/OS Communications Server: SNA Customization.

```
%!PS-Adobe-2.0
    614 25 translate 90 rotate .88 .76 scale /n 1 def /fs 10 def /ls 11. 2 def /ld l
```

Figure 55 on page 431 is a sample of LPR trace output for the following command with the XLATE option:

```
LPR test (p TIANNA h MVSA trace xlate GXS
```

In this sample, the server was not running, so the connection was not established. For detailed information about using and creating your own translate tables, see z/OS Communications Server: SNA Customization

```
EZB0915I Begin "LPR" to printer "TIANNA" at host "MVSA"
 1
EZB1057I Loaded translation table from "TCPUSR4.GXS.TCPXLBIN" .
EZB0920I Requesting TCP/IP service at 96155 20:04:14
EZB0921I Granted TCP/IP service at 96155 20:04:15
 2
EZB0922I Resolving MVSA at 96155 20:04:15
 3
EZB0924I Host MVSA name resolved to 9.67.113.60 at 96155 20:0 4:17
EZB0925I TCP/IP turned on.
EZB0926I Host "MVSA" Domain "TCP.RALEIGH.IBM.COM" TCPIP Service Machine TCP31S
EZB0927I Trying to open with local port 721 to foreign host address 9.67.113.60
 4
EZB1051E Failed to Open connection to Port Number = 515.  Return
             Code = -1.   Error Number = 61.  Port Number = 721.
             Remote IP Addr = 9.67.113.60
```

Figure 55. Example of LPR trace with XLATE option

Following are short descriptions of the numbered items in the trace:

1 Indicates the name of the translation table used by the LPR client. To avoid problems such as errors and data corruption, be sure that the LPD server is using the equivalent code pages.

2 Indicates the time the LPR client started trying to resolve the specified host name. The LPR client checks the name server table, the site, and address information files to resolve the host name.

3 Indicates the amount of time the LPR client took to resolve the specified host name. To reduce the amount of time, use the host IP address instead of the host name.

4 Indicates that the connection was not established. (In this sample, the LPD server was not running.) For a list of error numbers and their definitions, see z/OS Communications Server: IP and SNA Codes.

LPD server traces

This topic includes information about activating LPD server traces. It also provides samples of LPD trace output with explanations of selected messages.

Step for activating server traces

You can activate the tracing facilities within the LPD server in any of the following ways:

- Include the TRACE parameter in the LPSERVE PROC statement in the LPD server cataloged procedure.

  Be sure that a slash (/) precedes the first parameter and that each parameter is separated by a blank. For example:

  ```
  //LPSERVE PROC MODULE='LPD',PARMS='/TRACE'
  ```

- Enter the SMSG *procname* command, where *procname* is the name of the procedure used to start the LPD server.
- Specify the DEBUG statement in the LPD configuration file, LPDDATA.

Step for creating server trace output

LPD server traces go to the SYSPRINT data set. You can also define a DD card in the LPD cataloged procedure to write output to another data set. This section contains some samples of LPD server trace output.

Figure 56 is a sample of an LPD trace invoked by specifying the DEBUG option in the LPD configuration file, LPDDATA.

```
EZB0832I
EZB0621I LPD starting with port 515
EZB0679I Allocated ObeyBlock at 00005B70
EZB0679I Allocated ObeyBlock at 00005B60
EZB0679I Allocated ObeyBlock at 00005B50
EZB0628I Allocated PrinterBlock at 000058C0
EZB0629I    prt1 added.
EZB0641I Service prt1 defined with address
EZB0628I Allocated PrinterBlock at 00005630
EZB0629I    PRT1 added.
EZB0641I Service PRT1 defined with address
 1
EZB0628I Allocated PrinterBlock at 000053A0
EZB0629I    TIANNA added.
EZB0641I Service TIANNA defined with address
EZB0628I Allocated PrinterBlock at 00005110
EZB0629I    PRT2 added.
EZB0641I Service PRT2 defined with address
EZB0628I Allocated PrinterBlock at 000B1D40
EZB0629I    njesoto added.
EZB0641I Service njesoto defined with address
EZB0628I Allocated PrinterBlock at 000B1AB0
EZB0629I    rda added.
EZB0686I Host "9.37.33.159" resolved to 9.37.33.159.  Printer name is "lpt1".
EZB0641I Service rda defined with address
EZB0628I Allocated PrinterBlock at 000B1820
EZB0629I    POST added.
```

Figure 56. Example of LPD trace specified with the DEBUG option

```
 2
EZB0686I Host "9.67.105.55" resolved to 9.67.105. 55.  Printer name is "LPT2".
 2
EZB0641I Service POST defined with address
EZB0697I    ...End of Printer chain...
EZB0626I Allocated ConnectionBlock at 00147E08
 3
EZB0627I Passive open on port 515
EZB0705I 06/03/96 18:49:15
EZB0834I Ready
 4
EZB0789I GetNextNote with ShouldWait of TRUE
:
 5
EZB0790I GetNextNote returns.  Connection 1 NotificationConnection state changed (8681)
 5
EZB0779I New connection state Open (8673) on connection 1 with reason OK.
 5
EZB0782I Connection open.  Reading command.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0711I New command 2 data "2".
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification FSend response (8692)
EZB0799I Reading additional data on 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
 6
EZB0754I New subcommand 3 operands "7434 dfA827MV SA".
EZB0723I Allocated StepBlock at 000B1320
```

```
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns. Connection 2 Notification Connection state changed (8681)
EZB0779I New connection state Trying to open (8676) on connection 2 with reason OK.
EZB0626I Allocated ConnectionBlock at 0015BE08
```

7

```
EZB0627I Passive open on port 515
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 2 Notification Connection state changed (8681)
EZB0779I New connection state Open (8673) on connection 2 with reason OK.
EZB0782I Connection open.  Reading command.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification FSend response (8692)
EZB0799I Reading additional data on 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 2 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 2
EZB0711I New command 4 data "4".
EZB0708I FSend of response sent
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 2 Notification FSend response (8692)
EZB0763I Closing connection 2
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 2 Notification Connection state changed (8681)
EZB0779I New connection state Receiving only (8674) on connection 2 with reason OK.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification FSend response (8692)
EZB0799I Reading additional data on 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
```

8

```
EZB0754I New subcommand 2 operands "153 cfA827MVS A".
EZB0723I Allocated StepBlock at 000B1168
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification FSend response (8692)
EZB0799I Reading additional data on 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 2 Notification Connection state changed (8681)
EZB0779I New connection state Connection closing (8670) on connection 2 with reason OK.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification FSend response (8692)
EZB0799I Reading additional data on 1
```

```
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Connection state changed (8681)
EZB0779I New connection state Sending only (8675) on connection 1 with reason OK.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification FSend response (8692)
EZB0763I Closing connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Connection state changed (8681)
EZB0779I New connection state Connection closing (8670) on connection 1 with reason OK.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Connection state changed (8681)
EZB0779I New connection state Nonexistent (8672) on connection 1 with reason OK.
EZB0772I End Connection 1 for OK.
 9
EZB0776I Released StepBlock at 000B1320
 9
EZB0719I Allocated JobBlock at 00147798
 9
EZB0723I Allocated StepBlock at 000B1320
 10
EZB0716I Job 827 received prt1 MVSA
 11
EZB0734I Job 827 added to work queue
 12
EZB0716I Job 827 scheduled prt1 MVSA
EZB0776I Released StepBlock at 000B1168
EZB0777I Released ConnectionBlock at 0014AE08
EZB0824I ProcessWork starting on job queue
 13
EZB0731I Work Queue start
 13
EZB0732I $          827 JOBstartPRINTING
EZB0733I    Work Queue end
EZB0825I    Job 827 for prt1 dispatched in state JOBstartPRINTING
EZB0716I Job 827 printing prt1 MVSA
EZB0827I ProcessWork end with queue
EZB0731I     Work Queue start
 14
EZB0732I $          827 JOBcontinuePRINTING
EZB0733I     Work Queue end
EZB0789I GetNextNote with ShouldWait of FALSE
EZB0824I ProcessWork starting on job queue
EZB0731I     Work Queue start
EZB0732I $        827 JOBcontinuePRINTING
EZB0733I     Work Queue end
EZB0825I     Job 827 for prt1 dispatched in state JOBcontinuePRINTING
       flpNewBlock: State first call IsAtEof FALSE
 15
       flpNewBlock: State build      IsAtEof FALSE
       flpNewBlock: State check last IsAtEof FALSE
⋮
       flpNewBlock: State check last IsAtEof FALSE
       flpNewBlock: State build      IsAtEof FALSE
⋮
EZB0825I    Job 827 for prt1 dispatched in state JOBcontinuePRINTING
⋮
       flpNewBlock: State build      IsAtEof  TRUE
       flpNewBlock: State check last IsAtEof  TRUE
EZB0827I ProcessWork end with queue
EZB0731I       Work Queue start
EZB0732I $        827 JOBcontinuePRINTING
EZB0733I       Work Queue end
EZB0789I GetNextNote with ShouldWait of FALSE
EZB0824I ProcessWork starting on job queue
```

```
EZB0731I       Work Queue start
EZB0732I $        827 JOBcontinuePRINTING
EZB0733I       Work Queue end
EZB0825I    Job 827 for prt1 dispatched in state JOBcontinuePRINTING
EZB0827I ProcessWork end with queue
EZB0731I       Work Queue start
EZB0732I $        827 JOBfinishPRINTING
EZB0733I       Work Queue end
EZB0789I GetNextNote with ShouldWait of FALSE
EZB0824I ProcessWork starting on job queue
EZB0731I       Work Queue start
 16
EZB0732I $        827 JOBfinishPRINTING
EZB0733I       Work Queue end
EZB0825I    Job 827 for prt1 dispatched in state JOBfinishPRINTING
 17
EZB0716I Job 827 sent prt1 MVSA
 17
EZB0769I Job 827 removed from work queue
EZB0751I Released StepBlock at 000B1320
 17
EZB0716I Job 827 purged prt1 MVSA
EZB0771I Released JobBlock at 00147798
 18
EZB0827I ProcessWork end with queue
EZB0731I       Work Queue start
EZB0733I       Work Queue end
 19
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 0 Notification Connection state changed (8681)
 20
EZB0779I New connection state Nonexistent (8672)  on connection 0 with reason OK.
 20
EZB0772I End Connection 0 for OK.
EZB0777I Released ConnectionBlock at 00147E08
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 2 Notification Connection state changed (8681)
EZB0779I New connection state Nonexistent (8672) on connection 2 with reason OK.
EZB0772I End Connection 2 for OK.
EZB0777I Released ConnectionBlock at 0014DE08
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 3 Notification Connection state changed (8681)
EZB0779I New connection state Trying to open (8676) on connection 3 with reason OK.
EZB0626I Allocated ConnectionBlock at 00147E08
EZB0627I Passive open on port 515
EZB0789I GetNextNote with ShouldWait of TRUE
:
 21
EZB0790I GetNextNote returns.  Connection -48 Notification Other external interrupt received (8688)
 21
EZB0622I Terminated by external interrupt
```

Following are short descriptions of the numbered items in the trace:

1 Indicates that a control block was allocated for each service defined in the LPD configuration file. TIANNA is the name of one of the local printers.

2 Indicates that the remote printer, LPT2, was defined in a SERVICE statement with the name POST. LPT2 has the IP address 9.67.105.55.

3 Indicates that the LPD server listened on port 515 and that port 515 was opened.

4 Indicates that the LPD server waited for work.

5 Indicates that a connection was opened for an incoming LPR client and that the LPD server was receiving a command from that client.

6 Indicates that a subcommand was received from an LPR client. The subcommand indicates LPD was receiving a data file named dfA827MVSA, containing 7434 bytes of data. For details on commands and subcommands, see RFC117.

7 Indicates that the LPD server had a passive open connection on the restricted LPD port, 515.

8 Indicates that the LPD server was receiving a control file named cfA827MVSA, containing 153 bytes of data.

Note: Data files use the naming convention of df*x*. Control files use the naming convention cf*x*.

9 Indicates the control blocks that were allocated and released as files were received and processed. Control blocks are used primarily by IBM support for debugging purposes, in coordination with dumps.

10 Indicates that all data files for a particular job were received.

Note: Job number 827 is a three-digit job number generated by the LPR client.

11 Indicates that job 827 was added to this print queue. The LPD server maintains a work queue of jobs.

12 Indicates that job 827 was scheduled to be spooled to the output queue.

13 Indicates that the LPD server was processing print jobs from the work queue, and started sending print data to the JES output queue. The message `JOBstartPRINTING` does not mean that the file is physically printing.

14 Indicates that data was being sent for output. Depending on the size of the file, you might see this status many times for a single job.

15 Indicates checking for the end of the file as it is being processed. The number of `IsAtEof` entries depends on the data and size of the file.

16 Indicates that all data was processed and placed in the output queue.

17 Indicates that job 827 was completely processed by the LPD server and removed from the print queue, prt1, on host MVSA. Temporary data sets and control blocks for this job were also erased or released.

18 Indicates that the LPD server completed the jobs in that queue and scans the work queue again.

19 Indicates that the LPD server was waiting for more work to do.

20 Indicates that the LPR-to-LPD connection was closed normally.

21 Indicates that someone stopped the LPD server normally.

Figure 57 on page 437 is a sample of LPD trace output showing that job 947 failed to print because the client passed a filter that was not supported by the LPD server. In cases such as these, you can lose printouts. In this case, the LPD trace showed why, but the LPR trace did not show an error. (See Figure 52 on page 428 for the corresponding LPR trace output.)

```
EZB0831I IBM MVS LPD Version V2R10 on 05/05/98 at 19:21:46
EZB0832I
EZB0621I LPD starting with port 515
:
EZB0628I Allocated PrinterBlock at 000053A0
EZB0629I    TIANNA added.
EZB0641I Service TIANNA defined with address
:
EZB0627I Passive open on port 515
EZB0705I 06/03/96 19:21:47
EZB0834I Ready
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.Connection 0 Notification Connection state changed(8681)
EZB0779I New connection state Trying to open (8676) on connection 0 with reason OK.
EZB0626I Allocated ConnectionBlock at 0014AE08
EZB0627I Passive open on port 515
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.Connection 0 Notification Connection state changed(8681)
EZB0779I New connection state Open (8673) on connection 0 with reason OK.
EZB0782I Connection open.  Reading command.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0627I Passive open on port 515
:
 1
EZB0754I New subcommand 3 operands "333819 dfA947MV SA".
:
 2
EZB0754I New subcommand 2 operands "122 cfA947MVSA" .
:
EZB0776I Released StepBlock at 000B1438
EZB0719I Allocated JobBlock at 00147798
EZB0723I Allocated StepBlock at 000B1438
 3
EZB0716I Job 947 received TIANNA MVSA
 3
EZB0734I Job 947 added to work queue
 3
EZB0716I Job 947 scheduled TIANNA MVSA
EZB0776I Released StepBlock at 000B1280
EZB0777I Released ConnectionBlock at 0014AE08
EZB0824I ProcessWork starting on job queue
EZB0731I       Work Queue start
EZB0732I $        947 JOBstartPRINTING
EZB0733I       Work Queue end
EZB0825I    Job 947 for TIANNA dispatched in state JOBstartPRINTING
EZB0716I Job 947 printing TIANNA MVSA
```

Figure 57. Example of an LPD server trace of a failing job

```
 4
EZB0801I Filter "x" not supported.  Job abandoned.
EZB0827I ProcessWork end with queue
EZB0731I       Work Queue start
EZB0732I $        947 JOBfinishPRINTING
EZB0733I       Work Queue end
EZB0789I GetNextNote with ShouldWait of FALSE
EZB0790I GetNextNote returns.Connection 0 Notification Connection state changed(8681)
EZB0779I New connection state Connection closing (8670) on connection 0 with reason OK.
EZB0824I ProcessWork starting on job queue
EZB0731I       Work Queue start
EZB0732I $        947 JOBfinishPRINTING
EZB0733I       Work Queue end
 5
EZB0825I    Job 947 for TIANNA dispatched in state  JOBfinishPRINTING
```

```
EZB0716I Job 947 sent TIANNA MVSA
 6
EZB0769I Job 947 removed from work queue
EZB0751I Released StepBlock at 000B1438
 7
EZB0716I Job 947 purged TIANNA MVSA
EZB0771I Released JobBlock at 00147798
EZB0827I ProcessWork end with queue
EZB0731I       Work Queue start
EZB0733I       Work Queue end
⋮
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection -48 Notification Other external interrupt received (8688)
EZB0622I Terminated by external interrupt
```

Following are short descriptions of the numbered items in the trace:

1 Indicates that the LPD server received a command indicating the byte size and name of a data file sent by an LPR client.

2 Indicates that the LPD server received a command indicating the byte size and name of a control file sent by an LPR client.

3 Indicates that print job 947 was received, placed in the print queue named TIANNA on host MVSA, and was scheduled to be processed.

4 Indicates that the LPD server did not support filter x and discarded the print job.

5 Indicates that the job was finished. The flag `JOBfinishPRINTING` indicates the job is to be removed from the work queue and purged.

6 Indicates that the job was removed from the work queue and that the control blocks were released.

7 Indicates that the job was purged.

Figure 58 on page 439 is a sample of an LPD trace output generated by specifying the DEBUG statement in the LPD configuration file (LPDDATA). This sample shows that an LPR client issued a request, through an LPD server, to a printer defined as a remote server. (The LPD server acted as an LPR client by sending the request to a remote server.) Since the remote server was not running, the print job was purged.

Initially, the LPR client was unaware that the server was not running because the LPD server correctly acknowledged receipt of the data files and control files. Furthermore, the LPR trace did not indicate any problems. However, if you specify the option FAILEDJOB MAIL on the SERVICE statement for the remote printer, notification is sent to the user ID of the LPR client. For notification to be sent, Simple Mail Transfer Protocol (SMTP) must be running.

Note: The FAILEDJOB DISCARD option is the default.

The LPR lpd.config (p SOTO h MVS7 command was used to generate the trace output. SOTO is the name of the printer specified on the SERVICE statement, and MVS7 is the host on which the LPD server is running.

```
1
EZB0831I IBM MVS LPD Version V2R10
 on 05/05/98 at 19 :50:58
EZB0832I
EZB0621I LPD starting with port 515
EZB0679I Allocated ObeyBlock at 00005B70
EZB0679I Allocated ObeyBlock at 00005B60
EZB0679I Allocated ObeyBlock at 00005B50
EZB0628I Allocated PrinterBlock at 000058C0
EZB0629I    prt1 added.
EZB0641I Service prt1 defined with address
EZB0628I Allocated PrinterBlock at 00005630
EZB0629I    PRT1 added.
EZB0641I Service PRT1 defined with address
EZB0628I Allocated PrinterBlock at 000053A0
EZB0629I    TIANNA added.
EZB0641I Service TIANNA defined with address
EZB0628I Allocated PrinterBlock at 00005110
EZB0629I    PRT2 added.
EZB0641I Service PRT2 defined with address
EZB0628I Allocated PrinterBlock at 000B1D40
EZB0629I    njesoto added.
EZB0641I Service njesoto defined with address
EZB0628I Allocated PrinterBlock at 000B1AB0
EZB0629I    SOTO added.
2
EZB0686I Host "9.37.34.39" resolved to 9.37.34.39.   Printer name is "lpt1".
3
EZB0641I Service SOTO defined with address
EZB0628I Allocated PrinterBlock at 000B1820
EZB0629I    POST added.
EZB0686I Host "9.67.105.55" resolved to 9.67.105.55.  Printer name is "LPT2".
EZB0641I Service POST defined with address
EZB0697I    ...End of Printer chain...
EZB0626I Allocated ConnectionBlock at 00147E08
EZB0627I Passive open on port 515
EZB0705I 06/05/96 19:50:00
EZB0834I Ready
EZB0789I GetNextNote with ShouldWait of TRUE
:
EZB0782I Connection open.  Reading command.
EZB0799I Reading additional data on 1
:
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
4
EZB0754I New subcommand 3 operands "14221 dfA502MVS 7".
EZB0723I Allocated StepBlock at 000B1438
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0789I GetNextNote with ShouldWait of TRUE
:
EZB0799I Reading additional data on 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 1
```

Figure 58. Example of an LPD server trace for a remote print request

```
5
  19:50:48 EZB0754I New subcommand 2 operands "134 cfA502MVS7" .
     19:50:48 EZB0723I Allocated StepBlock at 000B1280
     19:50:49 EZB0789I GetNextNote with ShouldWait of TRUE
:
6
```

```
EZB0763I Closing connection 1
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Connection state changed (8681)
EZB0779I New connection state Connection closing (8670) on connection 1 with reason OK.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 1 Notification Connection state changed (8681)
EZB0779I New connection state Nonexistent (8672) on connection 1 with reason OK.
EZB0772I End Connection 1 for OK.
EZB0719I Allocated JobBlock at 00147798
 7
EZB0716I Job 502 received SOTO MVS7
 7
EZB0734I Job 502 added to work queue
 7
EZB0716I Job 502 scheduled SOTO MVS7
EZB0777I Released ConnectionBlock at 0014AE08
EZB0824I ProcessWork starting on job queue
EZB0731I Work Queue start
 8
EZB0732I $502 JOBstartSENDING
EZB0733I Work Queue end
EZB0825I Job 502 for SOTO dispatched in state JOBstartSENDING
EZB0626I Allocated ConnectionBlock at 0014AE08
 9
EZB0820I Trying to open with local port 721
 9
EZB0716I Job 502 opening SOTO MVS7
 10
EZB0769I Job 502 removed from work queue
EZB0827I ProcessWork end with queue
EZB0731I Work Queue start
EZB0733I Work Queue end
EZB0789I GetNextNote with ShouldWait of TRUE

.
.
.
EZB0790I GetNextNote returns.Connection 1 Notification Connection state changed(8681)
 11
EZB0779I New connection state Nonexistent(8672) on connection 1 with reason
      Foreign host did not respond within OPEN
 11
EZB0772I End Connection 1 for Foreign host
did not respond within OPEN timeout (8560).
EZB0705I 06/05/96 19:52:22
 12
EZB0773I Connection 1 terminated because "Foreign host did not respond within OPEN timeout (8560)"
 13 EZB0744I 748656 HELO MVS7.tcp.raleigh.ibm.com
 13 EZB0744I 748656 MAIL FROM:<LPDSRV3@MVSA>
 13 EZB0744I 748656 RCPT TO:<TCPUSR4@MVS7.tcp.raleigh.  ibm.com>
 13 EZB0744I 748656 DATA
 13 EZB0744I 748656 To:<TCPUSR4@MVS7.tcp.raleigh.ibm.com>
 13 EZB0744I 748656
 13 EZB0744I 748656 Your job to print the files "TCPUSR 4.LPD.CONFIG" on SOTO at MVSA has failed for
 13 EZB0744I 748656 this reason:  Remote connection
terminated (Foreign host did not respond within
 13 EZB0744I 748656 OPEN timeout (8560)).
 13 EZB0744I 748656 .
EZB0751I Released StepBlock at 000B1438
EZB0751I Released StepBlock at 000B1280
 14 EZB0716I Job 502 purged SOTO MVS7
EZB0771I Released JobBlock at 00147798
EZB0777I Released ConnectionBlock at 0014AE08
EZB0789I GetNextNote with ShouldWait of TRUE
 15 EZB0790I GetNextNote returns.  Connection 2 Notification Connection state changed (8681)
EZB0779I New connection state Nonexistent (8672) on connection2 with reason OK.
EZB0772I End Connection 2 for OK.
EZB0777I Released ConnectionBlock at 0014DE08
EZB0789I GetNextNote with ShouldWait of TRUE
```

```
EZB0790I GetNextNote returns.  Connection 0 Notification Connection
state changed (8681)
EZB0779I New connection state Trying to open (8676) on connection0 with reason OK.
EZB0626I Allocated ConnectionBlock at 00147E08
EZB0627I Passive open on port 515
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 0 Notification Connection state changed (8681)
EZB0779I New connection state Open (8673) on connection 0 with reason OK.
EZB0782I Connection open.  Reading command.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 0 Notification Data delivered (8682)
EZB0767I Timer cleared for connection 0
EZB0711I New command 4 data "4".
EZB0708I FSend of response sent
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 0 Notification FSend response (8692)
EZB0763I Closing connection 0
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection 0 Notification Connection state changed (8681)
EZB0779I New connection state Receiving only (8674) on connection0 with reason OK.
EZB0789I GetNextNote with ShouldWait of TRUE
EZB0790I GetNextNote returns.  Connection -48 Notification Other external interrupt received (8688)
EZB0622I Terminated by external interrupt
```

Following are short descriptions of the numbered items in the trace:

1 Indicates the date and time the LPD server was activated. This information can be compared to the date and time on an LPR trace to assure that both traces were generated for the same incident.

2 Indicates the IP address of the host. If the name of the host was specified instead of the IP address, this message would indicate if the IP address of the host was resolved.

3 Indicates that the name SOTO was defined on the SERVICE statement for the remote printer, lpt1, which had the address 9.37.34.39.

4 Indicates the byte size and the name of the data file sent from the LPR client on host MVS7.

5 Indicates the byte size and name of control file sent from the LPR client on host MVS7.

6 Indicates that the connection between the LPR client and the LPD server was closing, after the server received the data and control files.

7 Indicates that the print job was received, placed in the LPD print queue, represented by SOTO, and scheduled to be sent to its destination.

8 Indicates that the LPD server started to send print job 502 to the remote server.

Tip: If the printer was local, rather than remote, the message would have read `502 JOBstartPRINTING`.

9 Indicates that the LPD server, acting as a client, was opening a connection to the remote printer using local port 721.

10 Indicates that the LPD server removed the job from its work queue.

11 Indicates that the connection to the remote server timed out.

12 Indicates that the remote server did not respond to the request to open.

13 Indicates that the FAILEDJOB MAIL option was defined under the

SERVICE statement and that SMTP was running. The text in these messages was sent to the user ID of the LPR client.

14 Indicates that the print job was completely purged.

15 Describes additional activity between the LPD server and other clients.

Chapter 14. Diagnosing File Transfer Protocol (FTP) problems

This topic describes how to diagnose problems with the z/OS Communications Server FTP server and FTP client. If, after reading this topic, you are unable to solve your problem and you need to call the IBM Software Support Center, see one or both of the following sections for the documentation you need to provide: "Documenting server problems" on page 472 and "Documenting FTP client problems" on page 499.

This topic assumes your security product is RACF. However, you can use any SAF-compliant security product.

FTP server

This topic contains the following subtopics:

- "Structural overview"
- "Definitions and setup" on page 444
- "Error exit codes" on page 444
- "Name considerations for z/OS UNIX FTP" on page 444
- "Translation and data conversion support" on page 446
- "DB2 query support" on page 447
- "JES support" on page 449
- "Common z/OS UNIX FTP problems" on page 451
- "Diagnosing FTP server problems with traces" on page 464
- "Documenting server problems" on page 472

Structural overview

The z/OS model for the FTP server includes a daemon process and a server process. The daemon process starts when you start your cataloged procedure (for example, START FTPD) and it listens for connection requests on a specific port. The port is the well-known port 21 unless otherwise specified. For methods of choosing a different port number, see the information about configuring ETC.SERVICES and configuring the FTPD cataloged procedure in z/OS Communications Server: IP Configuration Guide. When the daemon accepts an incoming connection, it creates a new process (server's address space) for the FTP server, which handles the connection for the rest of the FTP login session. Each login session has its own server process.

The server process inherits the accepted connection from the daemon process. This connection is called the control connection. The server receives commands from the client and sends replies to the client using the control connection. The control connection port is the same as the daemon's listening port.

The client and server use a different connection for transferring data; this connection is called the data connection. By default, the data port is one less than the control connection port. For example, if the control connection port is 21, the data port is 20. An FTP client can override the default data port by directing the server to run in passive mode. In passive mode, the server uses an ephemeral port for the data port. Passive mode is requested by firewall-friendly clients and by clients initiating three-way data transfers.

© Copyright IBM Corp. 2000, 2015

Definitions and setup

This topic describes the definitions and setup for the FTP server.

Start procedure

The sample start procedure for the FTP server is EZAFTPAP (alias FTPD) in the SEZAINST data set. Changes might be necessary to customize the start procedure for your MVS host system.

Keep in mind the following for the FTP server start procedure:

- The library containing FTPD and FTPDNS must be APF authorized and must be either in the MVS link list or included on the STEPLIB DD statement.
- The C runtime libraries are needed for FTPD and FTPDNS. They must be APF authorized. If the C runtime library is not in the MVS link list, it must be included on the STEPLIB DD statement.
- If the FTP server is used for SQL queries, the DB2® DSNLOAD library must be APF authorized and must be either in the MVS link list or included on the STEPLIB DD statement.
- Several start options are available for the FTP server. If specified in the start procedure, these values override the default values for the FTP server and any values specified in the FTP.DATA data set.

For more information about the FTP server start procedure, see z/OS Communications Server: IP Configuration Reference.

FTP.DATA data set

The FTP.DATA data set is an optional data set that allows the FTP server configuration parameters to be customized. See z/OS Communications Server: IP Configuration Reference for more information about the FTP.DATA data set.

TCPIP.DATA data set

The TCPIP.DATA data set provides the following information to the FTP server:

- High-level qualifier to be used for configuration data sets
- Whether messages are to be written in uppercase or mixed-case
- Which DBCS translation tables are to be used

For more information about the TCPIP.DATA data set, see z/OS Communications Server: IP Configuration Reference.

Error exit codes

z/OS UNIX FTP uses the following error exit codes:

12 Daemon initialization failed; unable to accept an incoming connection. An EZY message identifying the specific problem is sent to syslogd.

24 The client session's initialization terminated because the FTP server load module cannot be loaded or executed. Message EZYFT53E is sent to syslogd.

28 Daemon initialization was terminated because the IBM TCP/IP is not enabled in the IFAPRD*xx* parmlib member. Message EZYFT54E is sent to syslogd and the operator console.

Name considerations for z/OS UNIX FTP

This topic explains the MVS and z/OS UNIX file system naming conventions.

MVS naming conventions

Restrictions: MVS data set names used with all FTP commands sent to the z/OS UNIX FTP server must meet MVS data set naming conventions as follows:

- Data set names cannot be longer than 44 characters.

  If the path name parameter sent with an FTP command is not enclosed in single quotation marks, the path name is appended to the current working directory to create the data set name. The combination of the current working directory and the path name cannot be longer than 44 characters. Issue the PWD command to display the current working directory.
- Each qualifier in a data set name, or each member name for a partitioned data set, must conform to the following:
  - No longer than 8 characters.
  - Begin with a letter or the special characters

    `$, @, or #.`
  - Contain only numbers, letters, or the special characters

    `$, @, #, -, or }.`
- Generation data group data set names must be in the format *gdg_name(generation_level)*. The *generation_level* is either 0, *+nn*, or *-nn*, where *nn* is the generation number. For example, the GDG data set MYGDG could be specified as `MYGDG(0)` for the current generation level, `MYGDG(-1)` for the next to the latest generation level, or `MYGDG(+1)` for the new generation level.

z/OS UNIX file system naming conventions

Guidelines: The following list describes some naming conventions you should know about when using z/OS UNIX file system files with the z/OS UNIX FTP server:

- The z/OS UNIX file system name is case-sensitive.
- If a name begins with a single quotation mark, specify `QUOTESOVERRIDE FALSE` in FTP.DATA, or use the SITE NOQUOTESOVERRIDE command.
- Names can contain imbedded blanks for special characters.

  Tip: Some FTP clients might truncate trailing blanks.
- The LIST and NLST subcommands, including all client subcommands that invoke the NLST subcommand, such as MGET or MDELETE, require special handling for certain special characters. For more information, see z/OS Communications Server: IP User's Guide and Commands.
- The START and SITE parameters have additional restrictions on the path name used with SBDATACONN. See z/OS Communications Server: IP Configuration Reference and z/OS Communications Server: IP User's Guide and Commands.
- When specifying a z/OS UNIX FTP subcommand with a file name containing special characters, some FTP clients might:
  - Truncate trailing blanks
  - Compress multiple internal blanks
  - Interpret special characters to have special meanings

  Unique specification of the file name such as enclosing in double or single quotation mark, or escaping special characters, might be necessary to make the client send the file name to the server correctly. See your client documentation to see whether this is necessary.

Translation and data conversion support

This topic describes translation and data conversion support for the FTP server.

Double-byte character set (DBCS) support

If you enter `quote type b <n>` at the client and if the DBCS translate table has not been loaded, the following reply is displayed:

```
504-Type not Supported. Translation table not loaded.
```

Do one or both of the following:

- Check the LOADDBCSTABLES statement in the TCPIP.DATA configuration file. If the statement wraps to the next line, parameters on the continued line are ignored. If all the parameters for the LOADDBCSTABLES statement do not fit on one line, use multiple LOADDBCSTABLES statements.
- Check the precedence order for TCPIP.DATA to ensure that the file being used contains the LOADDBCSTABLES statement or statements. Be aware that the location of TCPIP.DATA statements can be influenced in multiple ways, for example, by a GlobalTCPIPData specification or the RESOLVER_CONFIG environment variable. See z/OS Communications Server: IP Configuration Reference for the TCPIP.DATA search order.

Single-byte character (SBCS) support

Data conversion occurs for single-byte data on the data connection when ENCODING=SBCS is in effect and the data type is ASCII. For more information, see the FTP.DATA statement ENCODING in the z/OS Communications Server: IP Configuration Reference and the SITE ENCODING command in the z/OS Communications Server: IP User's Guide and Commands.

If you choose SBDATACONN as a statement in the FTP.DATA file or with the SITE SBDATACONN command, the FTP server builds a translation table using the code pages specified by SBDATACONN. If you receive the following reply to the SITE command, ask for a trace of the server with the UTL option to determine which characters cannot be translated.

200 Some characters cannot be translated between *codepage_1* and *codepage_2* .

If none of the untranslatable characters appear in the data, the data transfers are not affected. If, however, one of the untranslatable characters does appear, the data transfer fails and the client receives the following reply:

557 Data contains codepoints that cannot be translated.

You can avoid the failure if you specify a substitution character to replace non-translatable characters. For details on how to ask for character substitution, see SBSUB and SBSUBCHAR as FTP.DATA statements in the z/OS Communications Server: IP Configuration Reference and as parameters on the SITE command in z/OS Communications Server: IP User's Guide and Commands. If substitution occurs during the transfer, the client receives the following reply:

250 One or more characters were substituted during the transfer.

When substitution occurs at the destination of a data transfer, a subsequent transfer of the resulting data does not produce an exact copy of the original. For example, if you put a file to the server and one or more characters are substituted, the untranslatable characters are overlaid in the server copy with the substitution character. You cannot restore the original file by getting it from the server.

Multibyte character set (MBCS) support

Data conversion occurs for multibyte data on the data connection when ENCODING=MBCS is in effect and the data type is ASCII. For more information, see the FTP.DATA statement ENCODING in z/OS Communications Server: IP Configuration Reference and the SITE ENCODING command in z/OS Communications Server: IP User's Guide and Commands.

If you choose ENCODING=MBCS, you must specify MBDATACONN with a statement in the FTP.DATA file or with the SITE MBDATACONN command to name the code pages for the multibyte data transfer. If you attempt an ASCII data transfer with ENCODING=MBCS and no MBDATACONN specified, the client receives the following reply:

504 Multibyte encoding set but code pages are not defined.

If the multibyte data that you transfer has codepoints that cannot be translated, the transfer fails and the client receives the following reply:

557 Data contains codepoints that cannot be translated.

You can determine which bytes of the data cannot be translated by repeating the transfer with the DUMP 42 extended trace option active at the server.

DB2 query support

This topic describes how to use FTP server DB2 query support and how to diagnose SQL problems.

Steps for using FTP server SQL support

Complete these steps so that you can use the FTP server to submit queries to the DB2 subsystem.

Procedure

Perform the following steps:

1. Start the DB2 subsystem.
2. BIND the DBRM called EZAFTPMQ. This must be done whenever the part EZAFTPMQ.CSQLMVS has been recompiled.

   The DBRM must be bound into the plan named EZAFTPMQ, unless the keyword DB2PLAN was used in your FTP.DATA file to specify a different plan name.

   If you are running multiple instances of the z/OS UNIX FTP server at different maintenance levels, you must use DB2PLAN in FTP.DATA for each server and specify unique plan names.
3. Grant execute privilege to the public for the plan created in the previous step.

Results

To submit a query to DB2 through the FTP server, issue the following commands as necessary:

- SITE FILETYPE=SQL
- SITE DB2=*db2name*, where *db2name* is the name of a DB2 subsystem at the host
- RETR *fname1fname2*, where *fname1* is a file at the host that contains a SQL SELECT statement

Symptoms of SQL problems

Table 27 and Table 28 on page 449 show some symptoms and possible causes of SQL problems. Table 27 shows problems that generate a reply beginning with 55x.

Table 27. SQL problems generating 55x replies (FTP Server)

Reply	Output file	Possible causes
Reply 551: Transfer aborted: SQL PREPARE/DESCRIBE failure	The output file contains the SQL code and error message returned by the DB2 subsystem.	• A syntax error in the SQL statement in the host file. • The time stamp in the load module is different from the BIND time stamp built from the DBRM (SQL code = -818). This occurs if a BIND was not done for the EZAFTPMQ DBRM that corresponds to the current load module, or if the server is not configured to use the correct DB2 plan name. If this is the problem, every SQL query submitted through the FTP server fails.
Reply 551: Transfer aborted: unsupported SQL statement	No output is sent from the host.	The file type is SQL, but the host file being retrieved does not contain an SQL SELECT statement.
Reply 551: Transfer aborted: attempt to connect to *db2name* failed (*code*)	No output is sent from the host.	• The site db2name specifies a nonexistent DB2 subsystem. • The DB2 subsystem has not been started.
Reply 551: Transfer aborted: SQL not available. Attempt to open plan <planname> failed (DB2_reason_code).	No output is sent from the host.	• BIND was not done for the specified plan. • BIND was done for plan name other than EZAFTPMQ, but FTP.DATA does not contain a DB2PLAN statement to specify this planname. • User does not have execute privilege for the DB2 plan being used by the FTP server.
Reply 550: SQL query not available. Cannot load CAF routines.	No output is sent from the host.	The DSNLOAD library is not in the link list or the FTP server STEPLIB.
Note: For more information about the messages, see z/OS Communications Server: IP and SNA Codes.		

Table 28 on page 449 shows other SQL problems.

Table 28. Other SQL problems (FTP Server)

Problem	Possible causes
Output file contains only the SQL SELECT statement.	• The file type is SEQ, rather than SQL. If the file type is SEQ, a retrieve is done, but the host file is just sent back to the client. The query is not submitted to the DB2 subsystem. • The SELECT is for a VIEW for which the user ID does not have DB2 select privilege. The DB2 subsystem returns an empty table.
Client closes the connection because server is not responding.	The processing time needed by DB2 and FTP or both for the SQL query has exceeded the client's time limit for send or receive. An FTP server trace with the options FSC and SQL indicates the amount of SQL activity through FTP and the approximate time when each query was processed.

JES support

This topic describes the procedures to follow when JES output is not found and when remote job submission functions fail.

JES output not found (zero spool files)

In some cases, the server is in JESINTERFACELEVEL=1 and FILETYPE=JES, and a job has been submitted, but the output of the job cannot be found. You get zero spool files from a DIR command.

Use the following checklist to investigate:

1. Is the job name correct? The job name must consist of the user ID followed by a single character.
2. Was the job output spooled to the hold queue? The server can retrieve only job output that is in the hold queue. For JES/3, output must be assigned to an output queue held for external writer.
3. Did you set SBSENDEOL to a value other than CRLF for your original outbound file transfer? If so, it is not be possible to restart the file transfer. You should send the entire file to the server again.

If JESINTERFACELEVEL=2, ensure the JESJOBNAME, JESSTATUS and JESOWNER filters are set correctly with the STAT command.

If the server is in JESINTERFACELEVEL=2 and FILETYPE=JES, and a job has been submitted, but the output of the job cannot be found (that is, you get zero spool files from a DIR command), check the 125 reply message to verify that the JESOWNER, JESJOBNAME, and JESSTATUS filters are set to values that apply to your job.

If the JESJOBNAME=USER1* and the job submitted was USER2A, use the SITE command to set the JES filter to the appropriate value to find the job requested. If the SITE command does not allow the end user to change the values of the three

JES filters, see z/OS Communications Server: IP User's Guide and Commands to determine whether the proper Security Access Facility resources allow changing of the JES filters for the user.

Remote job submission functions fail

For problems with remote job submission, run the FTP JES trace to check for the following:

- Cannot allocate internal storage
- JES is not communicating
- JES unable to find output for the specified job ID
- Unable to acquire JES access
- Unknown return code from GET JES spool request
- JES unable to provide spool data set name now
- JES unable to get a job ID for a PUT or GET request
- JES PUT or GET aborted, job not found
- JES PUT or GET aborted, internal error
- JES PUT or GET aborted, timeout exceeded
- JES internal reader allocation failed
- JES user exit error

To trace the FTP JES activity, use the DEBUG=(JES) or DUMP=(JES) options of FTP syslog tracing. See "Diagnosing FTP server problems with traces" on page 464 for information about activating FTP syslog tracing.

FTP connection stops during FILETYPE=JES processing

This problem occurs due to extended periods of inactivity on the control or data connection while the FTP client is waiting for the FTP server to complete the job. Significant delays are possible during FILETYPE=JES processing due to conditions such as heavy system utilization and dispatching priorities.

You can avoid timeouts by using keepalive packets on the control and data connection. You can activate keepalive packets by:

- Coding the INTERVAL parameter on the TCPCONFIG statement in the TCP/IP profile (PROFILE.TCPIP). See z/OS Communications Server: IP Configuration Reference for more information about the TCPCONFIG statement.
- Coding the FTPKEEPALIVE statement in FTP.DATA. This activates keep alive packets for the control connection only, and overrides the TCPCONFIG INTERVAL statement. See z/OS Communications Server: IP Configuration Reference for more information about the FTPKEEPALIVE statement.
- Configuring the DATAKEEPALIVE option to activate keepalive packets for the data connection only. This overrides the TCPCONFIG INTERVAL statement in the stack. You can configure DATAKEEPALIVE by coding a statement in FTP.DATA, by using the **site** subcommand from the z/OS FTP client, or by sending a SITE command with the DATAKEEPALIVE parameter to the FTP server. See z/OS Communications Server: IP Configuration Reference for more information about the DATAKEEPALIVE statement.

Logging FTP server activity

The z/OS FTP server provides a way to log standardized information for the following types of activity:

- Connections from the client end user to the server

- Authentication of the client/server session (for example, through the use of Transport Layer Security)
- Access to the FTP server through User ID/password verification
- Allocation of MVS data sets and z/OS UNIX file system files
- Deallocation of MVS data sets and z/OS UNIX file system files
- Data transfers
- JES job submissions
- SQL queries
- Abnormal end (ABEND) conditions
- Confidence levels assigned to file transfers when CHKCONFIDENCE TRUE has been coded in FTP.DATA

Set the following server's FTP.DATA statements to enable logging:

```
FTPLOGGING
ANONYMOUSFTPLOGGING
```

For more information about these statements, see z/OS Communications Server: IP Configuration Reference.

Until the client sends the USER command to the server, the server cannot know whether this is an anonymous login. Therefore, up to the point the server processes the USER command, the FTPLOGGING statement and ANONYMOUSFTPLOGGING statement produce identical results.

This information is recorded in the SYSLOGD file. The data has an identification field that allows correlation of all entries for a specific login session.

For more information about configuring the SYSLOGD file, see z/OS Communications Server: IP Configuration Guide.

See z/OS Communications Server: IP Configuration Reference for the server's FTP.DATA configuration.

Common z/OS UNIX FTP problems

This topic describes some common z/OS UNIX FTP problems.

FTP daemon initialization problems

You might encounter the following problems when the FTP daemon is initialized.

No "Initialization Complete" message:
If the `EZY2702I Server-FTP Initialization completed at ...` message does not appear on the system console within a few minutes after starting the FTP daemon, verify that the daemon background job is still running. For example, if you started FTP with a procedure called FTPD, you can use the D A,L command to see whether the job FTPD1 is active.

If the background daemon job is running (for example, FTPD1), verify that TCP/IP is running. If it is not, start TCP/IP. The FTP initialization completes when TCP/IP starts.

If the background daemon job is not running, check the system console for nonzero exit codes from the background job. Look for messages in message or

trace output from syslogd for an EZY error message from FTP. The following are possible exit codes and the appropriate responses:

- 0012

  FTP is unable to use the port specified for the control connection. Look in the syslogd messages for the specific reason. Possible errors include the following:

  – `EZYFT13E bind error...Operation not permitted`

    Ensure that FTP has BPX.DAEMON authority.

  – `EZYFT13E bind error...Address already in use`

    Ensure that FTP is trying to use the correct port. The FTP server trace with the INT option indicates the port the daemon expects to use. If this is the correct port, you can use the TSO NETSTAT CONN command to determine the job that is currently using that port.

  – `EZYFT13E bind error...Permission denied`

    Ensure that the port you want FTP to use has been reserved for the FTP background job name. For example, if your start procedure is called FTPD and you want FTP to use port 21, the PORT statement in your *hlq*.PROFILE.TCPIP data set must specify `21 TCP FTPD1`.

- 0028

  This FTP daemon is not available because the IBM TCP/IP is not enabled.

Incorrect configuration values:
If you experience incorrect configuration values, check the following:

- Look in the syslogd output for message EZY2640I to verify that configuration values are coming from your intended FTP.DATA file. Verify that no errors were encountered reading this file.
- Determine whether your FTP.DATA file has sequence numbers. If it does, any statement with an optional parameter omitted picks up the sequence number as the parameter value.

  For example, the BLKSIZE statement has an optional parameter size. If you specify the size, the sequence number is ignored. If you do not specify the size, the system assumes that the sequence number is the size, causing an error.

FTP daemon not listening on expected port:
If the daemon is not listening on the expected port, verify that the correct port number is specified. Following is the preference order for a port number:

1. PORT start parameter
2. /etc/services
3. *hlq*.ETC SERVICES
4. A default port number of 21

AUTOLOG does not start the FTP daemon:
If your start procedure name contains fewer than 8 characters, ensure that the AUTOLOG and PORT statements in the *hlq*.PROFILE.TCPIP data set specify the FTP background job name. For example, if your start procedure is called FTPD, your *hlq*.PROFILE.TCPIP data set should specify FTPD1, as shown in the following examples:

- ```
  AUTOLOG
    FTPD JOBNAME FTPD1
  ENDAUTOLOG
  ```
- ```
  PORT
    20 TCP OMVS NOAUTOLOG   ;FTP data port
    21 TCP FTPD1            ;FTP control port
  ```

User exit routine is not invoked

If the user exit routine is not invoked, check the FTP trace in syslogd to see whether the exit routine was loaded. FTCHKIP is loaded once by the FTP daemon during initialization. The remaining user exits (FTCHKPWD, FTCHKCMD, FTCHKJES, FTPOSTPR, and FTPSMFEX) are loaded in the FTP server address space for each client session.

For example, check for one of the following responses:

```
main: ret code from fndmembr() for FTCHKIP is: 4
main: user exit FTCHKIP not found. Bypassing fetch().
```

or

```
main: ret code from fndmembr() for FTCHKCMD is: 0
main: chkcmdexit successfully loaded
```

If you have user-written exit routines and the FTP server is not able to find them, ensure that the user-written exit routines exist in an APF-authorized partitioned data set which is in the search order.

FTP Messages and FTP trace entries

If messages and trace entries do not appear in the syslog output file, do one or more of the following:

- Ensure that syslogd is configured for daemon entries. The file /etc/syslog.conf must have an entry for daemon.info to get FTP messages or an entry for daemon.debug to get FTP messages and trace entries.
- Ensure that the files specified for daemon entries exist at the time that syslogd started. If not, you need to create the files and recycle syslogd.
- Ensure that the files specified for daemon entries have appropriate permission bits (for example, 666).
- Ensure that syslogd is active.

If messages and trace entries display on the system console, it means that syslogd cannot write to the files specified for daemon entries and that /dev/console is defined. Check that syslogd is configured correctly and that the files specified for daemon entries have appropriate permission bits (for example, 666).

If you consider the volume of EZYFT47I messages logged by the server during initialization to be excessive, you can suppress these messages by adding a SUPPRESSIGNOREWARNINGS statement to the server's FTP.DATA. However, if you use this statement, the FTP server does not warn you when it ignores statements coded in FTP.DATA.

Guideline: Add SUPPRESSIGNOREWARNINGS to FTP.DATA only after you have verified all statements in FTP.DATA are correct.

FTP server abends

If the FTP server abends, check the following:

- S683 or U4088 abend validating user ID or password.
  - Ensure that the sticky bit has been turned on for the files /usr/sbin/ftpd and /usr/sbin/ftpdns.
  - Ensure that the FTPD and FTPDNS modules reside in an APF authorized partitioned data set, which is specified in the MVS linklist.
  - Ensure that all programs loaded into the FTP address space are APF authorized and are marked as controlled. This means that any FTP user exits,

the SQL load library, and the loaded run-time library need to be marked as controlled, using the RACF RDEFINE command. For more information, see z/OS UNIX System Services Planning, or see the RACF publications.

FTP session problems

The following topics describe some common FTP session problems.

Connection terminated by the server after user enters user ID:
The system console might display one of the following nonzero exit codes from the FTP server address space:

0012 This exit code indicates a socket error. See the syslogd messages for the specific error.

0024 This exit code indicates that the system was unable to load the server load module /usr/sbin/ftpdns. Ensure that the symbolic link or links for ftpdns are correct, that ftpdns exists in the z/OS UNIX file system and that the sticky bit is on, and that FTPDNS exists in the search order.

If your system is not configured to display exit codes, check the syslogd output for an FTP error message.

Connection terminated by the server after user enters password:
If the server terminates a connection after the user enters a password, ensure that the FTP load modules (FTPD and FTPDNS) reside in the APF authorized data set. Also, check that all programs accessed by the FTP address space are APF authorized and marked as "controlled." Additional symptoms include the following:

- The FTP daemon is running, but the FTP server address space abends.
- The FTP server trace is active with the ACC option, and the last FTP trace entry reads:

  `RA0nnn pass: termid is ...`

Connection terminated by the server after user enters any subcommand:
If the server terminates a connection after the user enters a subcommand, either one or both of the following events might occur:

- FTP server address space shows an exit code of `0000`.
- Last FTP server trace entry for the client session is `RXnnnn Server thread terminates rc = -2.` The preceding entries indicate a "select" error due to a bad file descriptor.

These events indicate that the server inactive time limit has probably expired with no activity from the client. If this happens frequently, check the inactive time set for the server. If necessary, increase it, and recycle the FTP daemon.

Password validation fails; session continues:
If password validation fails and the session continues, you receive the following reply:

```
530 PASS command failed
```

Additional replies might be generated if ACCESSERRORMSGS TRUE is coded in FTP.DATA.

If you receive this reply, do one or more of the following:

- Ensure all libraries, possibly indicated by ICH420I message, used by FTP are controlled and APF authorized.

- Ensure FTP is authorized if you are using BPX.DAEMON.
- Ensure that the FTP daemon has been started from a user ID running with superuser authority if the daemon has been started from the z/OS UNIX shell.
- Ensure that an OMVS segment is defined for the login user ID, or that automatic assignment of unique UNIX identities support is enabled.
- Obtain additional information about the error by enabling tracing with the ACC option.
- Ensure that the password or password phrase used to log in to FTP was entered correctly and is not expired.
- If you have activated mixed-case passwords in RACF or in another SAF-compliant security product, verify that you entered the password in the correct case. Password phrases used to log in to a z/OS FTP server are always case-sensitive even when mixed-case passwords are not enabled in your security product.
- If your password is comprised of more than one token, verify that your FTP client supports multiple token passwords. The z/OS FTP client supports multiple token passwords if you enclose the password with quotation marks when you code it in the NETRC file or data set, or enter it in response to a prompt.

Anonymous login fails:
If an anonymous login fails, use the following checklist to investigate:

1. Ensure that you have specified ANONYMOUS as a start parameter or in FTP.DATA.
2. Check the setting of the ANONYMOUSLEVEL variable in FTP.DATA. If ANONYMOUSLEVEL is not explicitly set in FTP.DATA, its value is equal to one.
3. If you have activated mixed-case passwords in RACF or in another SAF-compliant security product, verify the following:
   - The anonymous password in FTP.DATA is coded in the correct case
   - The anonymous password passed to the FTP daemon by the FTPD start procedure is coded in the correct case
   - The anonymous password specified by the MVS operator to override the parameters specified in the FTPD start procedure was coded in the correct case.

     Rule: Enclose the FTP parameters in single quotes when overriding the parameters specified in the FTPD start procedure while mixed-case passwords are enabled.

If ANONYMOUS is set in FTP.DATA, and the STARTDIRECTORY is in the z/OS UNIX file system, and ANONYMOUSLEVEL is two or three, verify that the required executable files are installed in the anonymous user's root directory. If the required executable files are not installed in the anonymous user's home directory, SYSLOGD contains error messages. For information about setting up the anonymous user's root directory, see z/OS Communications Server: IP Configuration Guide.

If you did not specify a user ID on the ANONYMOUS start parameter or on the FTP.DATA statement, ensure that the user ID ANONYMO is defined to TSO and RACF, and that it has a defined OMVS segment or that automatic assignment of

unique UNIX identities support is enabled. For information about the z/OS UNIX environment and its security considerations, see z/OS UNIX System Services Planning.

If you did specify a user ID on the ANONYMOUS start parameter or on the FTP.DATA statement, ensure that the specified user ID is defined to TSO and RACF, and that the specified user ID has a defined OMVS segment or that automatic assignment of unique UNIX identities support is enabled.

If ANONYMOUSLEVEL is two or three, verify that the STARTDIRECTORY value is compatible with the ANONYMOUSFILEACCESS value and that the FILETYPE value is compatible with the ANONYMOUSFILETYPESEQ, ANONYMOUSFILETYPEJES, and ANONYMOUSFILETYPESQL values.

If ANONYMOUSLEVEL=3 and if ANONYMOUS or ANONYMOUS/USERID/PASSWORD is coded, the user is prompted to enter an e-mail address as a password. Verify that the e-mail address entered by the user is consistent with the requirements of the EMAILADDRCHECK statement in FTP.DATA. If ANONYMOUS/USERID is coded, the user must provide the password for USERID. See z/OS Communications Server: IP Configuration Reference for more information about these FTP.DATA statements.

Wrong initial working directory:
If the initial working directory is *userid* instead of a z/OS UNIX file system directory, ensure that the STARTDIRECTORY statement is specified in the FTP.DATA data set and that the $HOME directory (defined or defaulted) exists for the login user ID.

Unable to open data connection message from server:
If, after issuing a command such as RETR, STOR, or LIST, the client receives the message `425 Unable to open data connection` from the server, check the FTP server trace for an error.

Tip: The trace option SOC should be active when you diagnose data connection errors.

See "Diagnosing FTP server problems with traces" on page 464 for information about starting the FTP server trace. One possible trace entry is `data_connect: bind() error...permission denied.` If you see this trace entry, ensure that the FTP data connection port is reserved to OMVS in the PROFILE.TCPIP data set.

```
PORT
  20 TCP OMVS NOAUTOLOG   ;FTP data port
  21 TCP FTPD1            ;FTP control port
```

Another possible trace entry is `data_connect: seteuid(0) error...Permission denied`. If you see this trace entry when the trace option ACC is active, ensure that FTP has BPX.DAEMON authority.

AT-TLS problems: The FTP server and client provide a level of security using the Application Transport Transparent Layer Security (AT-TLS) protocol. The FTP server and client use the services of System SSL as described in z/OS Cryptographic Services System SSL Programming. This document describes how system SSL works and also contains a topic about obtaining diagnostic information.

If you are experiencing problems with the AT-TLS support, gather AT-TLS trace information from FTP by activating security processing trace. You activate the trace before the FTP server starts by adding the DEBUG SEC statement to the server's FTP.DATA file or after the server starts (and before client connection) by using the MODIFY *jobname*,DEBUG=(SEC) operator command.

One of the common problems with the AT-TLS handshake is a mismatch in the ciphersuites supported by client and server. For a list of ciphersuites supported by z/OS FTP, see z/OS Communications Server: IP Configuration Reference.

```
SSL_NULL_MD5      ; 01
SSL_NULL_SHA      ; 02
SSL_RC4_MD5_EX    ; 03
SSL_RC4_MD5       ; 04
SSL_RC4_SHA       ; 05
SSL_RC2_MD5_EX    ; 06
SSL_DES_SHA       ; 09
SSL_3DES_SHA      ; 0A
```

Tip: Each ciphersuite has an associated number that is known to AT-TLS.

The following is a portion of the FTP server trace for a successful AT-TLS negotiation. In this example, the server of the FTP.DATA file was coded to accept only ciphersuites (cipherspecs) 01 and 02:

```
auth: entered with mechname TLS
ftpAuth: keyring = /u/user33/keyring/key.kdb
ftpAuth: stash   = /u/user33/keyring/key.sth
ftpAuth: environment_open()
ftpAuth: connect as a server
ftpAuth: environment_init()
ftpAuth: environment initialization complete
authClient: secure_socket_open()
authClient: cipherspecs = 0102
authClient: secure_socket_init()
tlsLevel: using TLSV1 with SSL_NULL_MD5 (01)
```

If the client were coded to not accept ciphersuites 01 and 02, the trace would look like this:

```
auth: entered with mechname TLS
ftpAuth: keyring = /u/user33/keyring/key.kdb
ftpAuth: stash   = /u/user33/keyring/key.sth
ftpAuth: environment_open()
ftpAuth: connect as a server
ftpAuth: environment_init()
tpAuth: environment initialization complete
uthClient: secure_socket_open()
uthClient: cipherspecs = 0102
uthClient: secure_socket_init()
uthClient: init failed with rc = 402 (GSK_ERR_NO_CIPHERS)
ndSecureConn: entered
EYFT96I TLS handshake failed
```

Data transfer problems

This topic describes various problems involving data transfer.

PASV and EPSV commands fail because no PASSIVEDATAPORTS are available:

If you code the PASSIVEDATAPORTS statement in the server's FTP.DATA, you must code enough ports to accommodate the server workload. Otherwise, EPSV and PASV commands to the server fail. Syslog tracing or CTRACE indicates bind() failed with errno 1116 - address not available, and errno2 of JRBINDNoPort.

To transfer data in passive mode, the FTP server must obtain a port from the PASSIVEDATAPORTS range. Therefore, allow at least one port per simultaneous data transfer. For example, if you expect one hundred users to log in to FTP at once to transfer data, code at least one hundred ports on the PASSIVEDATAPORTS statement.

The PASSIVEDATAPORTS statement does not preclude other applications from obtaining ports in the coded range. To prevent other applications from consuming ports in the PASSIVEDATAPORT range to the exclusion of FTP, code a PORTRANGE statement in PROFILE.TCPIP with the AUTHPORT parameter, specifying some or all ports in the PASSIVEDATAPORTS range. See z/OS Communications Server: IP Configuration Reference for more information about the PORTRANGE statement and PROFILE.TCPIP.

TCP/IP does not release ports that the FTP server has released until the connection associated with the port has exited the TIMEWAIT state. If all the PASSIVEDATAPORTS connections are in TIMEWAIT state, the server is not able to obtain a port to process a PASV or EPSV command. You can verify the connections are in TIMEWAIT state by issuing the **netstat -a** command from the USS shell. To correct this problem, increase the number of ports coded on the PASSIVEDATAPORTS statement .

Load module transfer failures:
This topic describes failures when transferring MVS load modules.

If the MVS load module transfers, but is not executable on the target system:

- Ensure that all hosts involved in the load module transfer are at the Communications Server for OS/390® V2R10 level or higher.
  - For proxy transfers, both servers and the client must be Communications Server for OS/390 V2R10 or higher.
- Ensure that the user did not attempt an operation that is not supported by load module transfer:
  - Ensure that the user did not attempt to rename the load module on transfer.
  - Ensure that the working directory on both the current and target systems is a load library of the correct type. An MVS load library for purposes of this support is a PDS with RECFM=U or a PDSE. Files can only be transferred between the same types of load libraries. This means that a PDS load library member must be transferred to another PDS, and a PDSE load library member must be transferred into another PDSE. The FTP client displays a terminal message `EZA2841I Local directory might be a load library` when a user changes local directory into a PDS or PDSE eligible for load module transfer support. The FTP server sends a `250-The working directory might be a load library` reply to the client when a CWD command is processed that causes the server working directory to become a PDS or PDSE eligible for load module transfer support. If both the message and the reply are not seen when changing directories before a transfer, load module transfer processing is not be used to transfer any files between the two directories.
  - Ensure that the load modules are transferred by member names only. The current working directory on both the target and destination systems must be the load library. Fully qualifying the member names is not permitted.
- Ensure that there are no problems with the IEBCOPY invocation. If an error is detected with an IEBCOPY invocation, the FTP server or client furnishes the IEBCOPY SYSPRINT output as messages to either the console (in the server's case) or the terminal session (in the client's case). Specify the FSC(2) debug

option for the general trace for the FTP client and for the FTP server to display the IEBCOPY SYSPRINT output for both successful and unsuccessful transfers. At the client, enter debug fsc(2) before the transfer. See "Start tracing" on page 465 for information about how to set the trace for the server.

If the MVS load module fails to transfer, check the following:

1. If `Reload of the load library failed` or `Unload of the load library failed` messages or replies are seen, then these messages indicate a problem with a call to the IEBCOPY system utility. Ensure that the IEBCOPY system utility is installed on the system and available to be called from application programs. If so, examine the FTP debug trace to determine whether IEBCOPY was successfully invoked (see the "Diagnosing FTP server problems with traces" on page 464 for information about activating FTP syslog tracing.) (Some client environments, particularly REXX scripts running under the UNIX System Services shell, are not fully authorized to call IEBCOPY). If IEBCOPY was successfully invoked, examine the IEBCOPY SYSPRINT output (described above) to see whether IEBCOPY reported any errors.
2. If allocation failure messages or replies are seen, then:
   - If the data set whose allocation failed is either the source or destination load library, ensure that no other process has allocated the load library for exclusive use.
   - If no data set name appears, or if the data set name ends in the characters XLMT, ensure that sufficient temporary DASD is available on the system. Load module transfer requires the use of sufficient temporary DASD to hold all data that could be transferred in one transfer command. Consider breaking up large mget or mput transfers into smaller groups to reduce the amount of required temporary DASD. If sufficient temporary DASD is not immediately available, then the setting of the AUTOMOUNT/NOAUTOMOUNT site option regulates whether FTP attempts to mount additional temporary storage to complete a load module temporary file allocation request.
3. If ABEND X'B3700000' occurs, this can indicate that a DASD storage shortage happened when FTP used IEBCOPY to copy the data from the original load module to a temporary data set. Ensure that all the members in the original load module PDS or library have correct size information in the directory. Do not transfer nonexecutable load modules, or load modules of size 0 or undefined size.

If the MVS load module transfer hangs, the system is probably waiting for temporary DASD to be mounted. If your system does not respond promptly to mount requests for temporary DASD, consider setting the NOAUTOMOUNT (LOC)SITE option about the hanging system, and breaking up large load module transfer mgets and mputs into smaller requests to reduce the requirement for temporary DASD.

Data set allocation fails:
If data set allocation is failing (MKD, STOR/STOU, or APPE), check for the following:

- Issue the STAT command and check for problems with the variables that define data set characteristics (LRECL, RECFM, BLKSIZE, PRIMARY, SECONDARY, or DIRECTORY).
  - Do they all have a valid value defined?
  - If the variable is not listed in the STAT command output, no value is assigned to this variable. If no value is assigned to the variable, the value must be

picked up from another source — either a model DCB or SMS. Does either the DCBDSN or DATACLASS (SMS) parameter have a valid value to provide a source for the missing variables?

- If an SMS data class is specified, is SMS active at the server system? (current SMS status is displayed as part of the output for the STAT command).
- If an SMS data class is specified, do the data class definitions contain values for the missing variables?
- Are both PRIMARY and SECONDARY either specified or not specified? If either PRIMARY or SECONDARY are specified, neither of the values is picked up from an SMS data class. Both must be unspecified to pick up the value from SMS or both must be specified to override the SMS values.
- If a model DCB is specified, are the characteristics of this data set valid for the data set being allocated?

- Issue the STAT command and check the PRIMARY, SECONDARY, and SPACETYPE values to determine how large the new data set is. The VOLUME and UNIT value of the STAT command indicate where the data sets are allocated. (If neither volume or unit is shown by the STAT command, data sets are allocated on the system default SYSDA DASD.) Does the server system have sufficient space where the data sets are allocated to allocate the data set? The SITE QDISK command provides information about the space available at the server system.
- Ensure that the destination at the server site is writable. Check with the operator at the server system to verify that the destination of the new data set is not write protected.

Data set allocation not picking up correct characteristics:
If the data set is being allocated successfully, but the resulting data set does not have the expected data set characteristics, check for the following:

1. All values obtained from SITE variables
   - Issue the STAT command to verify that the settings of all the SITE variables are correct. If any variables are missing from the STAT output, check for values specified for the DCBDSN or DATACLASS parameters. If a value is specified for the DCBDSN data set, go to Step 3 on page 461. If a value is specified for the DATACLASS parameter, go to Step 2.
   - Check for variables overridden by a client. The VM and MVS FTP clients automatically issue SITE commands when doing a STOR, STOU, or APPE command. The values sent automatically by the client could be overriding values set by specific SITE commands issued by the user. To prevent the VM or MVS client from automatically sending new SITE settings, issue the SENDSITE command at the client.
2. Values from SMS

   If the DATACLASS parameter has been specified, but the actual data set characteristics do not match the values in the specified SMS data class, issue the STAT command and check the information shown in the output from the STAT command for the following:

   - Is SMS active at the server system? If SMS is not active, the SMS data class cannot be used to define the data set.
   - Are values specified for any of the data set characteristic variables (DSNTYPE, LRECL, RECFM, BLKSIZE, PDSTYPE, PRIMARY, SECONDARY, RETPD, or DIRECTORY)? If these keywords are missing from the STAT output, or if the value 'SYSTEM' is displayed, no value is assigned to them and the data set characteristics should be picked up from the SMS data class. If, however, a value is present for any of these variables, the setting shown

by the STAT command overrides any information in the SMS data class. To pick up the value from the data class, issue the SITE command with the keyword with no value (for example, SITE RECFM) to turn off the parameter setting.

- Is a value specified for the DCBDSN parameter? If a DCBDSN data set is specified, the values for LRECL, RECFM, BLKSIZE, and RETPD are obtained from the model DCB data set and overrides any values in the SMS data class. Issue the SITE DCBDSN command to turn off the DCBDSN parameter setting.
- Check for variables overridden by a client. The VM and MVS FTP clients automatically issue SITE commands when doing a STOR, STOU, or APPE command. The values sent automatically by the client could be overriding values set by specific SITE commands issued by the user. To prevent the MVS or VM client from automatically sending new SITE settings, issue the SENDSITE command at the client.

3. Values from DCBDSN

   If the DCBDSN parameter has been specified, but the actual data set characteristics do not match the characteristics of the specified data set, issue the STAT command, and check the information shown in the output from the STAT command the following:

   - Are values specified for any of the data set characteristic variables (LRECL, RECFM, BLKSIZE, or RETPD)? If these keywords are missing from the STAT output, no value is assigned to them and the data set characteristics are picked up from the DCBDSN data set. If, however, a value is present for any of these variables, the setting shown by the STAT command overrides the values of the DCBDSN data set. To pick up the value from the DCBDSN data set, issue the SITE command with the keyword with no value (for example, SITE RECFM) to turn off the parameter setting.
   - Are variables being overridden by a client? The VM and MVS FTP clients automatically issue SITE commands when doing a STOR, STOU, or APPE command. The values sent automatically by the client could be overriding values set by specific SITE commands issued by the user. To prevent the VM or MVS client from automatically sending new SITE settings, issue the SENDSITE command at the client.

MVS data set not found:
If the server is not able to find the MVS data set, check for the following problems:

- Can the server find the data set to list it? Issue the DIR command to display the data set.
- Is the MVS data set at the server in the catalog? The server can only locate cataloged MVS data sets. Check the user level of access to the catalog. FTP servers at the z/OS V1R2 level and later display only the data sets to which the user has access.
- Was the *pathname* on the FTP command entered in single quotation marks? If not, the path name specified is appended to the end of the current working directory. Issue the PWD command to display the current working directory. If current_working_directory.pathname is not the correct name of the file, either change the current working directory with the CWD command or issue the correct data set name in single quotation marks as the *pathname*.

RETR, STOR, RNFR, RNTO, APPE, or DELE of data set fails:
If RETR, STOR, RNFR, RNTO, APPE, or DELE for the data set fails, check for the following problems:

1. Is the data set protected by a security system, such as RACF or permission bits or a retention period?
2. Is the data set being used at the server site by another program or user?
3. Was the data set available to the system, or was it migrated or on an unmounted volume?
4. Did the data set or member exist?
5. For RETR or STOR commands, did a REST command immediately precede the RETR or STOR?

   If so, the client is attempting to restart a file transfer. The server cannot detect certain REST argument errors until the RETR or STOR command is processed. If the trace options CMD and FSC are active, the server reply and server trace output provide insight into whether the REST command is implicated. Verify that the client and server have reestablished the original file transfer environment before attempting the restart.

The following problems apply to MVS data sets only:

1. Did the specified path name follow MVS data set naming conventions?
2. Was the requested data set a type of supported data set organization (PS, PDS, or PDS member) on a supported device type (DASD or tape)?
3. Were the path name specifications consistent with the type of data set? For example, if a member was requested, was the data set a PDS?

REST fails:
Use the STAT command to determine the current mode.

If mode is Block, report the problem to IBM.

If the mode is Stream, check the following:

- Verify that the server is configured for stream mode restarts. The server FEAT reply includes REST STREAM if the server is configured correctly.
- Inspect the REST reply for more insight into the reason the server rejected the REST command. See z/OS Communications Server: IP and SNA Codes for more information about FTP server replies.

Data transfer terminated: If data transfer terminated, check for the following problems:

1. Is the data set at the server large enough to receive the data being sent? If not, use the SITE command to change the space allocation for new data sets.
2. If storing a member of a PDS, is there room in the PDS for an additional member? Is there room in the PDS directory for another directory entry?
3. Did the client send an ABOR command?
4. Is the file type correct? For example, if filetype=SQL when it should be set to SEQ or JES, the host file being retrieved is assumed to be a SQL statement and FTP attempts to connect to DB2 and submit the statement to DB2 for processing.

Client abends during RETR command data transfer:
If the client abends while processing a RETR command, issue the STAT command, and check the value of the checkpoint interval. If this value is greater than zero and data is being transferred in EBCDIC, either block mode or compressed mode, the server is sending checkpoint markers with the data being transferred. If the client being used does not support checkpoint/restart, this checkpoint information

can cause unpredictable results, such as abends or data errors at the client. Change the setting of the checkpoint interval by issuing SITE CHKPTINT=0.

Data set disposition incorrect when transfer fails:
If the data set disposition is incorrect when transfer fails, check for the following problems:

- Data sets cataloged instead of deleted
  - Issue the STAT command and check the setting of the conditional disposition. If the STAT command output indicates `New data sets will be catalogued if a store operation ends abnormally`, the server catalogs new data sets, even if the data transfer fails. To change this setting, issue the SITE CONDDISP=DELETE command.
  - Did the transfer fail because the FTP server was either abending or being terminated by a STOP or CANCEL command? If this is the case, the data set is kept.
  - Is the client sending checkpoint information? If the data is being transferred in EBCDIC, either in block mode or compressed mode and the client has sent at least one checkpoint marker, the FTP server keeps the data set even if the conditional disposition is set to delete.
- Data sets deleted instead of cataloged
  - Issue the STAT command and check the setting of the conditional disposition. If the STAT command output indicates `New data sets will be deleted if a store operation ends abnormally`, the server deletes new data sets if the data transfer fails. To change this setting, issue the SITE CONDDISP=CATALOG command.

Checkpoint markers do not appear to be sent:
Issue the STAT command and check the settings for data transfer. Checkpoint information is only transferred in EBCDIC, with either block or compressed mode. The checkpoint interval must be greater than zero.

The sender of the data initiates the checkpoint information. Therefore, checkpointing must be set on at the client for a STOR, STOU, or APPE, (for the MVS FTP client, this is done by issuing the LOCSITE CHKPTINT=nn command with a value larger than zero) and set on at the server (by issuing the SITE CHKPTINT=nn command with a value larger than zero) for a RETR.

LOADLIB directory information is not sent with module transfer:
Issue the STAT command and check the settings for data transfer. Load module directory information is only sent for EBCDIC with a mode of either block or compressed.

Restriction: The client you are using must support the SDIR command.

Server PDS member statistics not created or updated:
ISPFStats must be set to TRUE to create or update the statistics for the PDS Member when using PUT, MPUT, GET, MGET, or APPEND subcommands. For PUT, MPUT, or APPEND, make sure the server's ISPFStats is set to TRUE. Issue the STAT command to determine this. If it is not set to TRUE, you can set it by using the SITE subcommand. For example, SITE ISPFStats sets ISPFStats to TRUE, and SITE NOISPFStats sets ISPFStats to FALSE.

Result: If the PDS directory block is full, PDS member statistics are not updated.

File transfers to the BatchPipe subsystem fail

If a file transfer to a batch pipe fails, the problem could be that the batch pipe reader has not been started. Verify that the batch pipe reader is active.

Guideline: In a JESMAS environment, if you use FTP to submit a job starting the batch pipe reader, the job can run on any system in the JESMAS environment unless you include this JCL statement in your job:

```
/*JOBPARM SASAFF=*
```

If your file transfer runs on a different system in the JESMAS environment, the JWT time limit will expire, and the FTP server JOB will appear to hang. You can avoid this by adding the JCL statement:

```
/*JOBPARM SASAFF=*
```

to your batch pipe job.

When a file transfer to a batch pipe fails with an allocation error, the problem could be that the BatchPipe subsystem has not been started. Verify that the BatchPipe subsystem is started. For example:

```
put 'user3.source.data' 'user3.subsys.output1'
 >>> SITE FIXrecfm 80 LRECL=80 RECFM=FB BLKSIZE=12960
 200 SITE command was accepted
 >>> PORT 9,42,104,22,4,15
 200 Port request OK.
 >>> STOR 'user3.subsys.output1'
 550 Allocation of USER3.SUBSYS.OUTPUT1 failed while executing STOR command.
 Command:
```

You can activate the FTP server trace to obtain more information about the failure. See “Diagnosing FTP server problems with traces” for information about activating the FTP server trace. Inspect the SYSLOG output for messages related to the file transfer failure.

A sample syslog output is:

```
Mar  2 16:29:47 MVS117 ftps[21]: GU1850 logALLOC: ALLOC  error in routine alloc_dasd
Mar  2 16:29:47 MVS117 ftps[21]: GU1852 logALLOC: SVC99 ALLOC failed with rc 4
Mar  2 16:29:47 MVS117 ftps[21]: GU1852 logALLOC: dsname   = USER3.SUBSYS.OUTPUT1
Mar  2 16:29:47 MVS117 ftps[21]: GU1852 logALLOC: S99ERROR = 04AC
Mar  2 16:29:47 MVS117 ftps[21]: GU2052 logDynErrMsg: entered
Mar  2 16:29:47 MVS117 ftps[21]: GU2088 logALLOC: ALLOC  Message=IKJ56231I
DATA SET USER3.SUBSYS.OUTPUT1 NOT ALLOCATED, SYSTEM OR INSTALLATION ERROR+
Mar  2 16:29:47 MVS117 ftps[21]: GU2088 logALLOC: ALLOC  Message=IKJ56231I
DYNAMIC ALLOCATION REASON CODE IS X'000004AC'
Mar  2 16:29:47 MVS117 ftps[21]: RS3033 mvs_store_data: allocation failed (4)

Allocation of USER3.SUBSYS.output1 failed (error code 04AC info code 0000 S99ERSN 00000000)
```

Error code 04AC means that the subsystem is not operational.

When FTPLOGGING is set to TRUE in the servers FTP.DATA file, the error messages will be logged in SYSLOGD. You can look up the messages and error codes.

Diagnosing FTP server problems with traces

Syslog tracing is available to aid in debugging z/OS UNIX FTP server problems. The following methods are available to start, stop, or modify syslog daemon and server tracing:

- TRACE start option
- FTP.DATA DEBUG statement
- FTP.DATA DUMP statement
- MODIFY *jobname*,DUMP operator command
- MODIFY *jobname*,DEBUG operator command
- Server SITE DEBUG command
- Server SITE DUMP command

See the following for more information:

- See "Start tracing during FTP initialization" and z/OS Communications Server: IP Configuration Reference, for details about the TRACE start option and FTP.DATA statements.
- See "Controlling the FTP server traces with MODIFY operator command" on page 467 and z/OS Communications Server: IP System Administrator's Commands for details about the MODIFY operator command.
- See "Stop tracing" on page 466, "Tracing activity for one user" on page 467, and the z/OS Communications Server: IP User's Guide and Commands for details about the SITE command.

After a client has logged in to FTP, the client can issue SITE DEBUG or SITE DUMP commands to change tracing for that session only.

Where to find traces

The z/OS UNIX FTP server sends its trace entries to syslogd. As shown in the following example, the daemon.debug statement in /etc/syslog.conf specifies where syslogd writes FTP trace records:

```
#
 # All ftp, rexecd, rshd
 # debug messages (and above
 # priority messages) go
 # to server.debug.a
 #
 daemon.debug              /tmp/syslogd/server.debug.a
```

All z/OS UNIX FTP trace entries are written to the same z/OS UNIX file system file.

Note: The TRACE parameter and MODIFY operator command options are issued to the FTP daemon and affect all client sessions that connect to the z/OS UNIX FTP server while tracing is active.

See the z/OS Communications Server: IP Configuration Guide for more information about syslogd.

Start tracing

This topic discusses the following methods of starting the FTP server traces:

- During FTP initialization
- After FTP initialization

Start tracing during FTP initialization:

You can use the TRACE start parameter, the TRACE statement, or the DEBUG and DUMP statements in FTP.DATA to begin tracing during FTP daemon initialization. This continues tracing for all FTP events for all FTP sessions.

Tracing remains active until you issue a MODIFY operator command to end it. See "Controlling the FTP server traces with MODIFY operator command" on page 467.

Tip: When you issue a MODIFY operator command to end tracing, tracing does not occur for any subsequent client sessions; however, tracing continues for any sessions that were already connected.

Start tracing after FTP initialization:
After initialization, you can enable tracing using an MVS MODIFY operator command to the FTP server listener process. See "Controlling the FTP server traces with MODIFY operator command" on page 467. Previously established FTP connections are not affected by a MODIFY operator command. Only FTP connections that are established after the MODIFY operator command was issued are subject to tracing.

If you have coded DEBUGONSITE TRUE and DUMPONSITE TRUE in the server's FTP.DATA file, you can use the SITE DEBUG command and the SITE DUMP command, respectively, to change tracing after you log in to FTP. For example, if you want to add JES general tracing and JES extended tracing, enter the following:

```
SITE DEBUG=(JES) DUMP=(JES)
```

If you want to restrict the use of the SITE command to change the tracing and your installation has a security product that supports the SERVAUTH class, you can provide additional levels of access control. If the installation has activated the SERVAUTH class and provided a profile for the SITE DEBUG command, only users who have read access to the profile are allowed to use the SITE DEBUG command. The profile name is:

```
EZB.FTP.systemname.ftpdaemonname.SITE.DEBUG
```

For example, if the procedure FTPD is used to start the server on system MVS164, the profile name is:

```
EZB.FTP.MVS164.FTPD1.SITE.DEBUG
```

The user's SITE DEBUG command is rejected if the security product determines that the user does not have read access to the profile.

If the installation has activated the SERVAUTH class and provided a profile for the SITE DUMP command, only users who have read access to the profile are allowed to use the SITE DUMP command. The profile name is:

```
EZB.FTP.systemname.ftpdaemonname.SITE.DUMP
```

For example, if the procedure FTPD is used to start the server on system MVS164, the profile name is:

```
EZB.FTP.MVS164.FTPD1.SITE.DUMP
```

The user's SITE DUMP command is rejected if the security product determines that the user does not have read access to the profile.

Stop tracing

Use the MODIFY operator command to stop global tracing. For example, your FTP jobname is FTPD1. You can issue the F FTPD1,DEBUG=(NONE) command to stop global tracing. Previously established FTP connections that were started with tracing enabled continue to produce trace output until the connections are terminated, but new connections start without tracing enabled.

If you have coded DEBUGONSITE TRUE in the server's FTP.DATA, the FTP client can use a SITE DEBUG=NONE command to stop tracing. The SITE command affects only tracing for the current FTP session.

Tracing activity for one user

A filter can be specified so that the traces are active only for certain clients that log in. Trace data can include both general and JES-related activity and includes data such as parameter lists and storage areas. The filtering can be done by either IP address of the client or by user ID for the session. Use the IPADDR(filter) and USERID(filter) operands on the FTP SITE command, or on MODIFY operator command, to enable trace filtering.

A client could use the SITE DEBUG and SITE DUMP subcommands to write excessive debugging information to the syslog and effectively disable the syslog function. To prevent this, a RACF profile controls whether a client is allowed to use these parameters on the SITE subcommand. FTP uses the SERVAUTH resource class. The resource name is `EZB.FTP.<systemname>.<ftpdaemonname>.SITE.<tracename>`. The lowest level is tracename, which is either DEBUG or DUMP.

Controlling the FTP server traces with MODIFY operator command

To start the general trace for the FTP server for all user IDs during initialization, specify the TRACE parameter either as a start option in the FTP server start procedure, or code a DEBUG BAS statement in FTP.DATA.

After initialization, use the MODIFY operator command to control the general and extended tracing for the FTP server. The command supports the following parameters:

- DEBUG for general tracing
- DUMP for extended tracing

Each allows a filter to be specified so that the traces are active for certain clients that log in. The filtering can be done by either IP address of the client or by user ID for the session.

Guideline: The *jobname* is the name associated with the FTP daemon background job. It is documented in message EZYFT41I in SYSLOGD. If you started the z/OS UNIX server using a procedure named FTPD, the job name to use for the MODIFY operator command is probably FTPD1. As client sessions connect to the FTP server, the session process adopts the trace options currently active. These options remain in effect for the life of the client session process, regardless of subsequent MODIFY operator commands issued to the FTP daemon.

Controlling general tracing:
To control the general trace, enter one of the following:

```
MODIFY jobname,DEBUG=(option_1,option_2,...,option_n,USERID(filter_name))

MODIFY jobname,DEBUG=(option_1,option_2,...,option_n,IPADDR(filter))
```

Where options are one of the following:

? Displays the status of the general traces.

The status of the trace is displayed as a response to all uses of the operator MODIFY DEBUG command. The ? allows you to get the status without making a change.

ACC
Shows the details of the login process.

ALL
Sets all of the trace points.

When the ALL parameter is processed, both the FSC and the SOC trace are set to level 1.

BAS
Sets a select group of traces that offer the best overall details without the copious output generated by certain trace options. Specifying this value is the same as the following:

```
MODIFY jobname,DEBUG=(CMD,INT,FSC,SOC)
```

CMD
Shows each command and the parsing of the parameters for the command.

FLO
Shows the flow of control within FTP. It is useful to show which services of FTP are used for an FTP request.

FSC*(n)*
Shows details of the processing the following file services commands

- APPE
- STOR
- STOU
- RETR
- DELE
- RNFR
- RNTO

This trace can be very intense; therefore, it allows you to specify levels of granularity for the trace points. The level 1 tracing that is specified by entering FSC or FSC(1) is the level normally used unless more data is requested by TCP/IP service group. The variable *n* can be a number in the range 1–8.

Level 1
Covers the major steps of the file services processing, which includes the following:

- Entry to a command processor
- Determination of the type of file being processed
- Choice of allocation method
- Choice of open method
- Choice of transfer routine
- Recognition of end of file or data
- Close and deallocation
- Call for SMF processing

Level 2
Provides more details for the major steps that are executed. These should be one-time events that enhance the information for the steps of level 1 tracing. An example would be some additional information about the allocation process.

Level 3
Provides trace information of repetitive events that occur during the

processing. For example, a trace for each full buffer (180K) of data that is received. Another example is a trace for each restart marker that is sent. The rate of repetition should be low enough that this level does not flood the trace.

Level 4
Provides trace information of repetitive events that occur at a higher rate than those of level 3. For example, a trace for each time data must be moved to the top of a buffer before the next receive_data.

Level 5
Provides trace information of repetitive events that occur at a higher rate than those of level 4. This is the most intense and covers events such as the processing of each block of data.

Tip: This level of tracing produces an extremely large amount of data and should not be used for large file transfers.

INT
The INT trace shows the details of the initialization and termination of the FTP session.

JES
The JES trace shows details of the processing for JES requests, such as when SITE FILETYPE=JES is in effect.

NONE
This value is used to turn off all of the traces.

PAR
The PAR trace shows details of the FTP command parser. It is useful for debugging problems in the handling of the command parameters.

SEC
The SEC trace shows the processing of security functions such as AT-TLS and GSSAPI negotiations.

SOC*(n)*
The SOC trace shows details of the processing during the setup of the interface between the FTP application and the network as well as details of the actual amounts of data that is processed. This trace can be very intense; therefore, it allows you to specify levels of granularity for the trace points. The level 1 tracing that is specified by entering SOC or SOC(1) is the level normally used unless more data is requested by the TCP/IP service group. The variable *n* can be a number from 1 to 8.

Level 1
Covers the major steps of the socket services processing. Connection initiation and closing steps are included.

Level 2
Adds more detail for level 1 events. For example, it traces the three steps that occur when a data connection is closed.

Level 3
The events for this trace are the send() and recv() calls for the data connection.

SQL
Shows details of the processing for SQL requests, such as when SITE FILETYPE=SQL is in effect.

UTL
Shows the processing of utility functions such as CD and SITE.

USERID(*filter_name*)
Filters the trace for user IDs matching the *filter_name* pattern.

If the user ID matches the filter at the time the clients log in, their tracing options are set to the current value of the options. Otherwise, tracing options are not set. Clients can use the SITE command to set their options after login if the initial ones are not appropriate. An example for the USERID filter is `MODIFY jobname,DEBUG=(CMD,USERID(USER3*))`, which activates the CMD trace for a user whose ID starts with `USER3`.

IPADDR(*filter*)
This optional parameter filters the trace for IP addresses matching the *filter* pattern.

If the IP address matches the filter at the time clients connect, its tracing options are set to the current value of the options. Otherwise, tracing options are not be set. Clients can use the SITE command to set their options after connect if the initial ones are not appropriate. An example of the IPADDR filter is `MODIFY jobname,DEBUG=(JES,IPADDR(9.67.113.57))`, which activates the JES trace for a client whose IP address is `9.67.113.57`. Another example is `MODIFY jobname,DEBUG=(JES,IPADDR(FEDC:BA98:7654:3210:FEDC:BA98:7654:3210))`. This activates the JES trace for a client whose IP address is FEDC:BA98:7654:3210:FEDC:BA98:7654:3210.

If the filter is an IPv4 address, submasking can be indicated by using a slash followed by a dotted decimal submask. For example, `192.48.32/255.255.255.0` allows addresses from `192.48.32.00` to `192.48.32.255`.

If the filter is an IPv6 address, network prefixing can be indicated by using a slash followed by a prefix length. For example, FEDC:BA98::0/32 allows all IP addresses from FEDC:BA98::0 to FEDC:BA98:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF.

The specification of the trace on the MODIFY operator command is not additive. That is, the trace setting is that of the last MODIFY operator command. For example:

```
MODIFY FTPDJG1,DEBUG=(NONE)
+EZYFT82I Active traces: NONE
MODIFY FTPDJG1,DEBUG=(CMD)
+EZYFT82I Active traces: CMD
MODIFY FTPDJG1,DEBUG=(FSC,USERID(USER33))
+EZYFT82I Active traces: FSC(1)
+EZYFT89I Userid filter: USER33
MODIFY FTPDJG1,DEBUG=(SOC)
+EZYFT82I Active traces: SOC(1)
```

Guidelines: The following are some guidelines to use for migrating from previous versions of the MODIFY operator command :

- `MODIFY jobname,TRACE`

  This is still accepted and is equivalent to `MODIFY jobname,DEBUG=(BAS)`. The old response message EZY2704I is replaced by EZYFT82I.
- `MODIFY jobname,NOTRACE`

  This is still accepted and is equivalent to `MODIFY jobname,DEBUG=(NONE)`. The old response message EZY2705I is replaced by EZYFT82I.
- `MODIFY jobname,JTRACE`

  This is still accepted and is equivalent to `MODIFY jobname,DEBUG=(CMD,FSC,JES)`. The old response message EZY2710I is replaced by EZYFT82I.

- `MODIFY jobname,NOJTRACE`

  This is still accepted and is equivalent to `MODIFY jobname,DEBUG=(NONE)`. The old response message EZY2711I is replaced by EZYFT82I.
- `MODIFY jobname,UTRACE=USER33`

  This is rejected as an obsolete command. Its function can be replaced with the following pair of commands:

  ```
  MODIFY jobname,DEBUG=(ALL,USERID(USER33))
  MODIFY jobname,DUMP=(ALL,USERID(USER33))
  ```
- The use of the ALL parameter can produce an extensive amount of trace data and should not be specified on a routine basis.
- `MODIFY jobname,NOUTRACE`

  This is rejected as an obsolete command. If complete tracing was activated as suggested in the previous step, then the tracing can be stopped as follows:

  ```
  MODIFY jobname,DEBUG=(NONE)
  MODIFY jobname,DUMP=(NONE)
  ```

Controlling extended tracing:
To control the extended trace, enter one of the following:

```
MODIFY jobname,DUMP=(option_1,option_2,...,option_n,USERID(filter_name))

MODIFY jobname,DUMP=(option_1,option_2,...,option_n,IPADDR(filter))
```

Where options are one of the following:

id Specifies the ID number of a specific extended trace point that is to be activated in the FTP code. The ID number has a range of 1–99.

? Displays the status of the extended traces.

ALL
Activates all of the trace points.

NONE
Resets (turns off) all extended traces.

FSC
Activates all of the extended trace points in the file services code. The numbers activated are 20–49.

SOC
Activates all of the extended trace points in the network services code. The numbers activated are 50–59.

JES
Activates all of the extended trace points in the JES services code. The numbers activated are 60–69.

SQL
Activates all of the extended trace points in the SQL services code. The numbers activated are 70–79.

USERID(*filter_name*)
Filters the trace for user IDs matching the *filter_name* pattern.

If a client's user ID matches the filter when the client logs into the server, its tracing options are set to the current value of the options. Otherwise, tracing options are not set. Clients can use the SITE command to set their options after login if the initial ones are not appropriate. An example for the USERID filter is `MODIFY jobname,DEBUG=(21,USERID(USER33))`, which activates the dumpID 21 trace for a user if his user ID is `USER33`.

IPADDR(*filter*)
Filters the extended trace for IP addresses matching the *filter* pattern.

If the client's IP address matches the filter when the client connects to the FTP server, its extended tracing options are set to the current value of the options. Otherwise, tracing options are not set. Clients can use the SITE command to set their options after connect if the initial ones are not appropriate.

An example of the IPADDR filter is `MODIFY jobname,DUMP=(JES,IPADDR(9.67.113.57))`, which activates the JES extended trace for a client whose IP address is `9.67.113.57`. Another example is `MODIFY jobname,DUMP=(FSC,IPADDR(FEDC:BA98:7654:3210:FEDC:BA98:7654:3210))`. This activates all file services extended traces for a client whose IP address is FEDC:BA98:7654:3210:FEDC:BA98:7654:3210.

If the filter is an IPv4 address, submasking can be indicated by using a slash followed by a dotted decimal submask. For example, `192.48.32/255.255.255.0` allows addresses from `192.48.32.00` to `192.48.32.255`.

If the filter is an IPv6 address, network prefixing can be indicated by using a slash followed by a prefix length. For example, FEDC:BA98::0/32 allows all IP addresses from FEDC:BA98::0 to FEDC:BA98:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF.

The specification of the trace on the MODIFY operator command is not additive. That is, the trace setting is that of the last MODIFY operator command. For example:

```
MODIFY FTPDJG1,DUMP=(NONE)
+EZYFT83I Active dumpIDs: NONE
MODIFY FTPDJG1,DUMP=(21)
+EZYFT83I Active dumpIDs: 21
MODIFY FTPDJG1,DUMP=(22)
+EZYFT83I Active dumpIDs: 22
```

Guidelines: The following are guidelines for migrating from the old parameters that were used with the MODIFY operator command:

- `MODIFY jobname,DUMP`

  This format is rejected. DUMP requires at least one parameter (see above).
- `MODIFY jobname,NODUMP`

  This is still accepted and is equivalent to `MODIFY jobname,DUMP=(NONE)`. The old response message EZY2656I is replaced by EZYFT83I.
- `MODIFY jobname,JDUMP`

  This is rejected as an obsolete command with a suggestion to use the DUMP parameter. For example, use the `MODIFY jobname,DUMP=(JES)` command.
- `MODIFY jobname,NOJDUMP`

  This is rejected as an obsolete command with a suggestion to use the DUMP parameter. For example, use the `MODIFY jobname,DUMP=(NONE)` command.

Documenting server problems

If the problem is not caused by any of the common errors described in this section, collect the following documentation before calling the IBM Support Center.

Documentation is divided into the following categories: essential and helpful (but not essential).

- Essential
  - Precise description of problem, including expected results and actual results

- z/OS UNIX FTP server dump (for abends)
- z/OS UNIX FTP server traces (see "Diagnosing FTP server problems with traces" on page 464 for information about collecting FTP server traces)
  - Minimum for initial problem reporting: DEBUG BAS

- Helpful
  - FTP client output
    - If the FTP client is a z/OS client, include a trace in the output by one of these methods:
      - Coding DEBUG statements in the client's FTP.DATA file.. See z/OS Communications Server: IP Configuration Reference for information about the DEBUG statement.
      - Invoking the FTP client with the -d or TRACE invocation option. See z/OS Communications Server: IP User's Guide and Commands for more information.
      - Specifying a DEBUG subcommand in the client command input stream before the affected transfer. Use this option only if the problem does not involve the initial establishment of the FTP control session.
  - Server FTP.DATA data set
  - TCPIP.DATA data set
  - PROFILE.TCPIP data set
  - ETC.SERVICES data set
  - The reply from the STAT or XSTA command issued to the server.

    Guidelines:
    - Issue the STATus subcommand from the z/OS FTP client to retrieve STAT command output from the server. From non z/OS clients, you may have to issue QUOTE STAT to retrieve the output from the server.
    - Issue the STATus subcommand with a parameter from the z/OS FTP client to issue the XSTA command to the server. The XSTA command retrieves output related to that parameter only. From z/OS clients prior to V1R8, or from non-z/OS clients, issue QUOTE XSTA *(parameter* to retrieve the output from the server.
  - Any console messages issued for resources experiencing errors.
  - If applicable, sample data to re-create the problem

FTP client

This topic describes the following subtopics:

- "Execution environments" on page 474
- "Setup" on page 474
- "Naming considerations" on page 475
- "Directing the client to exit when an error occurs" on page 475
- "Translation and data conversion support" on page 475
- "File tagging support" on page 476
- "DB2 query support" on page 480
- "Restarting file transfers" on page 482
- "Diagnosing FTP connection and transfer failures with EZA2589E" on page 483
- "Problems starting the client" on page 491
- "Problems logging in to the server" on page 492

- "Problems transferring data" on page 494
- "Other problems" on page 497
- "Diagnosing FTP client problems with tracing" on page 497
- "Documenting FTP client problems" on page 499

Execution environments

The FTP client can run in any of the following environments:

- Interactive (under the TSO or the z/OS UNIX shell)
- Batch (under TSO only)
- REXX exec (under TSO)

When run interactively, you can redirect terminal I/O. When run under TSO, server responses and debug messages can be redirected to a file. For example, you can use the **ftp 9.68.100.23 > 'USER27.FTPOUT'** command to redirect output from a TSO command line to a data set. When run under the z/OS UNIX shell, both input and output can be redirected. To redirect input from the file /user27/ftpin and output to the file /user27/ftpout, issue the **ftp 9.68.100.23 > /user27/ftpout < /user27/ftpin** command.

Tip: When redirecting output under z/OS UNIX, nothing is displayed on the system console, not even command prompts, and it is difficult to know when input is requested. Consequently, use output redirection only when also using input redirection.

Setup

Use an FTP.DATA data set to customize configuration parameters. You can use a SOCKSCONFIGFILE data set or file to instruct the client to connect to certain FTP servers through a SOCKS server. For information about the FTP.DATA data set and SOCKS configuration data set or file used by the FTP client, see z/OS Communications Server: IP User's Guide and Commands, z/OS Communications Server: IP Configuration Guide, and z/OS Communications Server: IP Configuration Reference.

Message EZY2640I displays the name of the FTP.DATA file. Use the FTP client **locstat** subcommand to display the name of the SOCKS configuration data set or file that is being used.

The TCPIP.DATA configuration file provides information for the FTP client, such as the high-level qualifier to be used for configuration data sets and which DBCS translation tables can be used. For more information about the TCPIP.DATA configuration file, see z/OS Communications Server: IP Configuration Reference.

Tip: The z/OS UNIX search order for the file is used even if the FTP client is invoked under TSO.

You can install the security exits for the FTP client to control FTP client commands or monitor replies from the FTP server. See z/OS Communications Server: IP Configuration Guide, z/OS Communications Server: IP Configuration Reference, and z/OS Communications Server: IP User's Guide and Commands for more information about the security user exits for the FTP client.

Naming considerations

The FTP client can access both MVS data sets and z/OS UNIX file system files. For more information, see "Name considerations for z/OS UNIX FTP" on page 444.

Directing the client to exit when an error occurs

You can direct the FTP client to exit whenever an error occurs, rather than to continue processing. You also have some control over whether the client exits with a generic return code or with a return code that reflects the type of error that occurred. For a description of all the FTP client return code options, see z/OS Communications Server: IP User's Guide and Commands.

Translation and data conversion support

This topic describes translation and data conversion support for the FTP client.

Double-byte character set (DBCS) support

If the DBCS translate tables are not available, the client issues the following message after a valid command to establish a double-byte transfer type (for example, SJISKANKI, BIG5, or 'TYPE B n') is entered:

```
"EZA1865I Command not Supported. Translation Table not Loaded.
```

If this message is displayed, check the LOADDBCSTABLES statement in the TCPIP.DATA file. If the statement wraps to the next line, parameters on the continued line are ignored, and no error message is issued. If all parameters for the LOADDBCSTABLES statement do not fit on one line, use multiple LOADDBCS statements.

Check the precedence order for the TCPIP.DATA file to ensure that the file being used contains the LOADDBCSTABLES statement or statements. Be aware that the location of TCPIP.DATA statements can be influenced in multiple ways. For example, by a GlobalTCPIPData specification or the RESOLVER_CONFIG environment variable. See z/OS Communications Server: IP Configuration Guide for the TCPIP.DATA search order.

Single-byte character (SBCS) support

Data conversion occurs for single-byte data on the data connection when ENCODING=SBCS is in effect and the data type is ASCII. For more information, see the FTP.DATA statement ENCODING and the LOCSITE ENCODING subcommand in the z/OS Communications Server: IP User's Guide and Commands.

If you choose SBDATACONN as a statement in the FTP.DATA file or with the LOCSITE SBDATACONN subcommand, the FTP client builds a translation table using the code pages specified by SBDATACONN. If you receive the following message from the LOCSITE subcommand, start the trace with the DEBUG UTL option to determine which characters cannot be translated:

EZYFS08I
Some characters cannot be translated between *codepage_1* and *codepage_2*.

If none of the untranslatable characters appear in your data, your data transfers are not affected. If an untranslatable character is present in the data you are trying to transfer, your data transfer fails and you receive the following message:

EZA2930I
Transfer failed because data cannot be translated.

To avoid the failure, specify a substitution character to replace non-translatable characters. For more information about how to specify character substitution, see SBSUB and SBSUBCHAR as FTP.DATA statements and as parameters on the LOCSITE subcommand in z/OS Communications Server: IP User's Guide and Commands. If substitution occurs during the transfer, you receive the following message:

EZA2947I
One or more characters were substituted during the transfer.

When substitution occurs at the destination of a data transfer, a subsequent transfer of the resulting data does not produce an exact copy of the original. For example, if you get a file from the server and one or more characters are substituted, the untranslatable characters are overlaid with the substitution character. You cannot restore the original file by putting it to the server.

Multibyte character set (MBCS) support

Data conversion occurs for multibyte data on the data connection when ENCODING=MBCS is in effect and the data type is ASCII. For more information, see the FTP.DATA statement ENCODING and the LOCSITE ENCODING subcommand in z/OS Communications Server: IP User's Guide and Commands.

If you choose ENCODING=MBCS, you must specify MBDATACONN with a statement in the FTP.DATA file or with the LOCSITE MBDATACONN subcommand to name the code pages for the multibyte data transfer. If you attempt an ASCII data transfer with ENCODING=MBCS and no MBDATACONN specified, you receive the following message:

EZZ9793I
Multibyte encoding requested but code pages are not defined.

If the multibyte data that you transfer has codepoints that cannot be translated, the data transfer fails and you receive the following message:

EZA2930I
Transfer failed because data cannot be translated

To determine which bytes of the data cannot be translated, repeat the transfer with the DUMP 42 extended trace option active at the client.

File tagging support

When the server writes a z/OS UNIX file system file, it might tag the file using the USS support for file tagging. In some cases you might experience conflicts when you try to read a file that has been tagged. A tagged file has a file tag, which is an attribute that identifies the coded character set ID (ccsid) of the text data within the file. When a tagged file is read from the file system, the data is translated using the ccsid if SBDATACONN has specified a network transfer code page to use with the file's code page. A file might also be untagged or tagged binary.

ASCII file transfers

If you put data into a z/OS UNIX file system file when the data type is ASCII, the file is tagged if you have used SBDATACONN to specify the code page for the file system and for the network transfer. That is, you have specified SBDATACONN=(file_system_cp,network_transfer_cp). If the data conversion table is the FTP_STANDARD_TABLES or is specified using XLATE, the file is not tagged. The following client session example shows the effects of combining data type ASCII and SBDATACONN defined tables using code pages:

```
. . .
 1 (01) Command:  ascii
(02) >>> TYPE A
(03) 200 Representation type is Ascii NonPrint
 2 (04) Command:  site sbd=(ISO8859-1,ISO8859-1)
(05) >>> SITE sbd=(ISO8859-1,ISO8859-1)
(06) 200 Site command was accepted
 3 (07) Command:  put a afile
(08) >>> PORT 9,67,113,57,4,121
(09) 200 Port request OK.
(10) >>> STOR afile
(11) 125 Storing data set /u/user33/tagging2/afile
(12) 250 Transfer completed successfully.
(13) 200 bytes transferred in 0.070 seconds.  Transfer rate 2.86 Kbytes/sec.
 4 (14) Command:  site sbd=(IBM-1047,ISO8859-1)
(15) >>> SITE sbd=(IBM-1047,ISO8859-1)
(16) 200 Site command was accepted
 5 (17) Command:  put a efile
(18) >>> PORT 9,67,113,57,4,122
(19) 200 Port request OK.
(20) >>> STOR efile
(21) 125 Storing data set /u/user33/tagging2/efile
(22) 250 Transfer completed successfully.
(23) 200 bytes transferred in 0.005 seconds.  Transfer rate 40.00 Kbytes/sec.
 6 (24) Command:  site sbd=FTP_STANDARD_TABLES
(25) >>> SITE sbd=FTP_STANDARD_TABLES
(26) 200 Site command was accepted
 7 (27) Command:  put a ufile
(28) >>> PORT 9,67,113,57,4,123
(29) 200 Port request OK.
(30) >>> STOR ufile
(31) 125 Storing data set /u/user33/tagging2/ufile
(32) 250 Transfer completed successfully.
(33) 200 bytes transferred in 0.005 seconds.  Transfer rate 40.00 Kbytes/sec.
 8 (34) Command:  ls -T
(35) >>> PORT 9,67,113,57,4,124
(36) 200 Port request OK.
(37) >>> NLST -T
(38) 125 List started OK
 9 (39) t ISO8859-1   T=on  afile
 10 (40) t IBM-1047    T=on  efile
 11 (41) - untagged      T=off ufile
(42) 250 List completed successfully.
 12 (43) Command:  get afile
(44) >>> PORT 9,67,113,57,4,125
(45) 200 Port request OK.
(46) >>> RETR afile
 13 (47) 125-Tagged ASCII file translated with current data connection translation table
(48) 125 Sending data set /u/user33/tagging2/afile
(49) 250 Transfer completed successfully.
(50) 190 bytes transferred in 0.005 seconds.  Transfer rate 38.00 Kbytes/sec.
(51) Command:  get efile
(52) >>> PORT 9,67,113,57,4,126
(53) 200 Port request OK.
(54) >>> RETR efile
 13 (55) 125-Tagged EBCDIC file translated with current data connection translation table
(56) 125 Sending data set /u/user33/tagging2/efile
(57) 250 Transfer completed successfully.
(58) 200 bytes transferred in 0.005 seconds.  Transfer rate 40.00 Kbytes/sec.
(59) Command:  get ufile
(60) >>> PORT 9,67,113,57,4,127
(61) 200 Port request OK.
(62) >>> RETR ufile
 14 (63) 125 Sending data set /u/user33/tagging2/ufile
(64) 250 Transfer completed successfully.
(65) 200 bytes transferred in 0.005 seconds.  Transfer rate 40.00 Kbytes/sec.
```

```
15 (66) Command:  site sbd=(IBM-1047,ISO8859-1)
(67) >>> SITE sbd=(IBM-1047,ISO8859-1)
(68) 200 Site command was accepted
(69) Command:  get afile
(70) >>> PORT 9,67,113,57,4,128
(71) 200 Port request OK.
(72) >>> RETR afile
16 (73) 125-Tagged ASCII file translated with table built using file system cp=ISO8859-1,
              network transfer cp=ISO8859-1
(74) 125 Sending data set /u/user33/tagging2/afile
(75) 250 Transfer completed successfully.
(76) 190 bytes transferred in 0.005 seconds.  Transfer rate 38.00 Kbytes/sec.
(77) Command:  get efile
(78) >>> PORT 9,67,113,57,4,129
(79) 200 Port request OK.
(80) >>> RETR efile
17 (81) 125-Tagged EBCDIC file translated with table built using file system cp=IBM-1047,
              network transfer cp=ISO8859-1
(82) 125 Sending data set /u/user33/tagging2/efile
(83) 250 Transfer completed successfully.
(84) 200 bytes transferred in 0.005 seconds.  Transfer rate 40.00 Kbytes/sec.
18 (85) Command:  ebcdic
(86) >>> TYPE E
(87) 200 Representation type is Ebcdic NonPrint
(88) Command:  get afile
(89) >>> PORT 9,67,113,57,4,142
(90) 200 Port request OK.
(91) >>> RETR afile
19 (92) 557 File contains ASCII data - enter TYPE A command before entering RETR command
(93) Command:  get efile
(94) >>> PORT 9,67,113,57,4,143
(95) 200 Port request OK.
(96) >>> RETR efile
20 (97) 125 Sending data set /u/user33/tagging2/efile
(98) 250 Transfer completed successfully.
(99) 190 bytes transferred in 0.005 seconds.  Transfer rate 38.00 Kbytes/sec.
```

Notes:

1. Change the data type to ASCII.
2. Site command requests a file system code page ISO8859-1, an ASCII code page.
3. Put a file and name it afile.
4. Site command requests a file system code page IBM-1047, an EBCDIC code page.
5. Put a file and name it efile.
6. Site command requests standard FTP translation tables.
7. Put a file and name it ufile.
8. Use the **ls** subcommand to determine whether files in a UNIX System Services file system are tagged (that is, have a file tag). You use the -T option to request the file tagging information. When options are specified on the **ls** subcommand, name parameters cannot be specified.
9. afile is a tagged file. Its file system code page is IS08859-1. It is a Text file.
10. efile is a tagged file. Its file system code page is IBM-1047. It is a Text file.
11. ufile is an untagged file. It is not a Text file.
12. Retrieve afile, which is a tagged file.

13 Client receives an indication that the tagged file is translated using the current tables because the current data connection tables were not specified with a network transfer code page (see **6**).

14 Since this is an untagged file, no indication is needed about the tables used.

15 Specify translation tables with a file system code page and a network transfer code page.

16 The code page of the tagged ASCII file is used with the network transfer code page to translate the data in the file.

17 The code page of the tagged EBCDIC file is used with the network transfer code page to translate the data in the file.

18 Change the data type to EBCDIC.

19 The 557 reply informs the client that the data type must be ASCII when the file that is tagged as ASCII is retrieved.

20 The EBCDIC file is OK to send with data type EBCDIC since no translation occurs and the data is already EBCDIC.

Binary file transfers

If you put data into a z/OS UNIX file system file when the data type is binary, the file is tagged as a binary file. The following client session example shows the effects of the binary file tagging:

```
. . .
 1 (01) Command:  binary
(02) >>> TYPE I
(03) 200 Representation type is Image
(04) Command:  put a file
(05) >>> PORT 9,67,113,57,4,44
(06) 200 Port request OK.
(07) >>> STOR file
(08) 125 Storing data set /u/user33/newtag/file
(09) 250 Transfer completed successfully.
(10) 190 bytes transferred in 0.050 seconds.  Transfer rate 3.80 Kbytes/sec.
(11) Command:  ascii
(12) >>> TYPE A
(13) 200 Representation type is Ascii NonPrint
 2 (14) Command:  ls -T
(15) >>> PORT 9,67,113,57,4,45
(16) 200 Port request OK.
(17) >>> NLST -T
(18) 125 List started OK
 3 (19) b binary      T=off file
(20) 250 List completed successfully.
(21) Command:  get file
(22) >>> PORT 9,67,113,57,4,46
(23) 200 Port request OK.
(24) >>> RETR file
 4 (25) 557 File contains binary data - enter TYPE I command before entering RETR command
(26) Command:  binary
(27) >>> TYPE I
(28) 200 Representation type is Image
(29) Command:  get file
(30) >>> PORT 9,67,113,57,4,47
(31) 200 Port request OK.
(32) >>> RETR file
(33) 125 Sending data set /u/user33/newtag/file
(34) 250 Transfer completed successfully.
(35) 190 bytes transferred in 0.005 seconds.  Transfer rate 38.00 Kbytes/sec.
```

Notes:

1. Request binary data type.
2. Use the **ls** subcommand to determine whether files in a UNIX System Services file system are tagged. You use the -T option to request the file tagging information. When options are specified on the **ls** subcommand, name parameters cannot be specified.
3. The tagging information shows that the file is a binary file.
4. The 557 reply informs the client that the data type must be binary when the file is retrieved.

DB2 query support

This topic describes how to use the FTP client DB2 query support and how to diagnose SQL problems.

Steps for using FTP client SQL support

Complete these steps so that you can use the FTP client to submit queries to the DB2 subsystem.

Procedure

Perform the following steps:

1. Start the DB2 subsystem.
2. BIND the DBRM called EZAFTPMQ. This must be done whenever the part EZAFTPMQ.CSQLMVS has been recompiled.

   The DBRM must be bound into the plan named EZAFTPMQ, unless the keyword DB2PLAN was used in your FTP.DATA file to specify a different plan name.
3. Grant execute privilege to the public for the plan created in the previous step.

Results

To use the FTP client to submit a query to DB2 and send the output to the FTP server, issue the following commands as necessary:

- LOCSITE FILETYPE=SQL
- LOCSITE DB2=*db2name* where *db2name* is the name of a DB2 subsystem at the local host
- PUT *fname1 fname2* where *fname1* is a local file that contains a SQL SELECT statement

Symptoms of SQL problems

Table 29 on page 481 and Table 30 on page 482 show some symptoms and possible causes of SQL problems.

Table 29 on page 481 shows problems that generate a reply beginning with 55x.

Table 29. SQL problems generating 55x replies (FTP Client)

Reply	Output file	Possible causes
EZA2570E: Transfer aborted: SQL PREPARE/DESCRIBE failure	The output file contains the SQL code and error message returned by the DB2 subsystem.	• A syntax error in the SQL statement in the host file. • The time stamp in the load module is different from the BIND time stamp built from the DBRM (SQL code = -818). This occurs if a BIND was not done for the EZAFTPMQ DBRM that corresponds to the current load module, or if the server is not configured to use the correct DB2 plan name. If this is the problem, every SQL query submitted through the FTP server fails.
EZA2573E: Transfer aborted: unsupported SQL statement	No output is sent from the host.	The file type is SQL, but the host file being retrieved does not contain an SQL SELECT statement.
EZA2568E: Transfer aborted: attempt to connect to *db2name* failed (*code*)	No output is sent from the host.	• The locsite db2name specifies a nonexistent DB2 subsystem. • The DB2 subsystem has not been started.
EZA2569E: Transfer aborted: SQL not available. Attempt to open plan <planname> failed (DB2_reason_code).	No output is sent from the host	• BIND was not done for the specified plan. • BIND was done for plan name other than EZAFTPMQ, but FTP.DATA does not contain a DB2PLAN statement to specify this plan name. • User does not have execute privilege for the DB2 plan being used by the FTP server.
EZA2740E: SQL query not available. Cannot load CAF routines.	No output is sent from the host.	The DSNLOAD library is not in the link list or the FTP server STEPLIB.
Note: For more information about these messages, see z/OS Communications Server: IP and SNA Codes.		

Table 30 on page 482 shows other SQL problems.

Table 30. Other SQL problems (FTP Client)

Problem	Possible causes
Output file contains only the SQL SELECT statement.	• The file type is SEQ, rather than SQL If the file type is SEQ, a retrieve is done, but the local file is sent to the server. The query is not submitted to the DB2 subsystem. • The SELECT is for a VIEW for which the user ID does not have DB2 select privilege. The DB2 subsystem returns an empty table.
Connection terminated.	The processing time needed by DB2 or FTP or both for the SQL query has exceeded the server time limit for send or receive. If you are using the MVS FTP server and the server trace shows a select error due to a bad file descriptor, check the inactive time set for the server and, if necessary, increase the time. An FTP client trace indicates the amount of SQL activity through FTP and the approximate time when each query is processed.

Restarting file transfers

A valid restart of an interrupted file transfer depends on reestablishing the environment that existed at the time the file transfer failed. Environment includes:

- The current FTP.DATA statements
- The current SITE and LOCSITE settings
- The sequence of commands (such as Type, Mode, and Structure) that affect the way FTP transfers files
- The current translation tables in use on the data connection

Restriction: All environment settings must be re-created before attempting to restart a file transfer.

The following sections describe some possible problems that you might encounter.

Client rejects the RESTART subcommand

Use the following checklist if the client rejects the RESTART subcommand:

- Verify that you have re-created the original file transfer environment.
- Verify that your environment met all the restrictions for the RESTART subcommand.
- Verify that checkpointing was active during the failed file transfer.
- See z/OS Communications Server: IP User's Guide and Commands for information about RESTART subcommand restrictions and checkpointing a file transfer

Client rejects SRESTART subcommand

Use the following checklist if the client rejects the SRESTART subcommand:

- Verify that you have re-created the original file transfer environment.

- Verify that the environment met the SRESTART subcommand restrictions.
- See z/OS Communications Server: IP User's Guide and Commands for information about SRESTART subcommand restrictions. Unlike the restart subcommand, you do not need to activate check pointing, but you do need to enter the SRESTART parameters correctly

Client accepts SRESTART subcommand, but server rejects RESTART

Use the following checklist if the client accepts SRESTART subcommand, but server rejects RESTART:

- Verify that the server supports stream mode restarts by issuing a FEAT command to the server. The FEAT reply l includes the keyword REST_STREAM if the server supports stream mode restarts.
- Some FTP servers other than z/OS FTP servers reply to the FEAT command with REST_STREAM when they support stream mode restarts in one direction only, such as server to client file transfers. Contact the provider of the FTP server software to verify the server support stream restarts for the direction of the transfer you are attempting.
- See z/OS Communications Server: IP User's Guide and Commands for information about the feature subcommand.
- Verify that the server has recreated the environment extant during the failed file transfer.
- Did you restart a retrieve and SBSENDEOL is configured to a value other than CRLF at the server, and the server is a z/OS FTP server? If so, you cannot restart the file transfer. Retrieve the file from the server again.

Diagnosing FTP connection and transfer failures with EZA2589E

EZA2589E is issued to describe a timeout or interruption while the FTP client was processing. The following example shows the message format:

```
EZA2589E Connection to server interrupted or timed out. operation
```

The message indicates the operation that was in progress when the FTP interruption occurred. Each operation is listed below along with the timer being used and the suggested response. Timers can be set individually in FTP.DATA, or all the timers can be set to one value using the (TI xx or -t xx option when starting the FTP client. See z/OS Communications Server: IP User's Guide and Commands for more information regarding the timers.

If the message was generated due to a user interruption, such as using Ctrl-C, ensure the FTP client had enough time to complete before being interrupted. In some cases, a packet trace or a CTRACE might be required to determine why the message was issued. See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45 or INFOAPAR II12014 for instructions for taking packet traces and CTRACES.

For more information about EZA2589E, see z/OS Communications Server: IP Messages Volume 1 (EZA).

Values and explanations for operation in EZA2589E

This topic lists and describes the values for operation in EZA2589E.

Initial Connection

Timer MYOPENTIME

Explanation
The FTP client is trying to establish a connection with the FTP server. Either the TCP connection has not completed yet or the initial reply from the server has not been received.

User Response
Ensure the remote server responds to a ping request. The value of MYOPENTIME can be increased to allow more time for the server to send the initial reply. If the problem recurs, contact the system programmer.

System Programmer Response
If there are firewalls between the FTP client and FTP server, ensure the firewalls are allowing FTP traffic from the client IP address to the FTP server for the port being used. A packet trace of the failing transfer shows whether the TCP connection has been completed, the IP addresses being used, and any replies sent by the server.

Initial IPv6 connection

Timer MYOPENTIME

Explanation
The FTP client is trying to establish a connection with an FTP server using an IPv6 address. Either the TCP connection has not completed yet, or the initial reply from the server has not been received.

User Response
Ensure the remote server responds to a ping request. The value of MYOPENTIME can be increased to allow more time for the server to send the initial reply. If the problem recurs, contact the system programmer.

System Programmer Response
If there are firewalls between the FTP client and FTP server, ensure the firewalls are allowing IPv6 FTP traffic from the client to the FTP server for the port being used. A packet trace of the failing transfer shows if the TCP connection has been completed, the IP addresses being used, and any replies sent by the server.

Waiting for data connection

Timer INACTTIME

Explanation
The FTP client is waiting for the FTP server to establish a data connection. A PORT or EPRT command, shown in a previous EZA1701I message, has been sent to the FTP server indicating the IP address and port on which the client is listening. The server should initiate a TCP connection to the FTP client. This connection has not completed yet.

User Response
Increase the value of INACTTIME and try again. Contact the system programmer if the failure recurs.

System Programmer Response
Ensure that active data connections or PORT or EPRT commands are allowed by any firewalls between the client and server. Take a packet trace of the failure to determine whether the remote FTP server has attempted the connection to the FTP client. If the packet

trace does not show an SYN packet arriving from the server to the specified IP address and port, investigate the FTP server and the path to the FTP client to determine whether the connection is being blocked. If the FTP client is not responding to the SYN packet, take a CTRACE (with options TCP and INTERNET) and a packet trace. Send these to IBM customer service. The FTP client could also be configured to use firewall friendly data connections by issuing the **locsite fwfriendly** subcommand before the **get** or **put** subcommand or by coding FWFRIENDLY TRUE in FTP.DATA. This might allow the data connection to complete because it causes the client to send a PASV or EPSV command instead of a PORT or EPRT command.

Guideline: The PORT or EPRT command sent to the server determines the port and IP address the FTP server connects to. For EPRT, the format is EPRT |X|Y|Z|, where X is the address family, Y is the IPv4 or IPv6 address and Z is the port. For the PORT command, the port being used must be calculated. For PORT, the format is a,b,c,d,x,y, where a.b.c.d is the IPv4 address, and (x * 256) + y is the port number.

Sending a command

Timer INACTTIME

Explanation
The FTP client has timed out sending a command to the FTP server. This indicates that the TCP layer is unable to transmit data to the remote server.

User Response
Contact the system programmer.

System Programmer Response
Take a packet trace to investigate the TCP traffic between the two hosts.

Sending ABORT command

Timer INACTTIME

Explanation
The FTP client has timed out sending an ABORT command to the FTP server. This indicates that the TCP layer is unable to transmit data to the remote server.

User Response
Contact the system programmer.

System Programmer Response
Take a packet trace to investigate the TCP traffic between the two hosts.

Receiving data

Timer DATACTTIME

Explanation
The FTP client is waiting for data from the FTP server on the data connection. A full buffer of data has not arrived within the DATACTTIME seconds, or the FTP client was interrupted by the user before a full buffer of data arrived. The FTP client issues a

recv() call, which returns only when its buffer is full or when the connection has ended. The FTP client uses a default buffer size of 180K. The FTP client is dependent on the data connection closing cleanly. This informs the FTP client that all the data has arrived from the server. If the connection does not close cleanly, this message is issued.

User Response

Increase the DATACTTIME to allow more time for data to arrive. If the failure recurs, contact the system programmer.

System Programmer Response

Take a packet trace to investigate the data transfer. The packet trace should be analyzed for conditions which would slow down the transfer, such as retransmitted packets or decreasing window sizes. Increasing the DATACTTIME can allow the FTP client more time to recover from these types of network issues. DATACTTIME should also be increased for transfers over low-bandwidth connections, such as dialup. If the packet trace shows that the connection does not close cleanly (for example, the FIN packet is not properly acknowledged), the remote server might need to be investigated as well.

Tip: For best results, specify the Session option when formatting the packet trace.

Sending data

Timer DATACTTIME

Explanation

The FTP client has timed out sending data to the FTP server over the data connection. The FTP client sets a timer to the value of DATACTTIME seconds before issuing a send call. If the send does not complete in that time period or the FTP client is interrupted by the user, the FTP transfer fails. This timeout can be caused by a slowdown in the transfer, such as network congestion or the remote machine not accepting data.

User Response

Increase the value of DATACTTIME to allow more time for the data transmission to occur. If the failure recurs, contact the system programmer.

System Programmer Response

Take a packet trace to investigate the data transfer. Analyze the trace for causes of a slowdown. For slow networks, such as dialup, increase the DATACTTIME. If the packet trace shows many retransmitted packets, investigate the network to determine why packets are being dropped.

The window size advertised by the FTP server can also slow down the connection. If the FTP server is advertising a small window size, investigate the server to determine whether the window size can be increased. If the FTP server is very busy, causing the window size to decrease or even go to 0, increase the DATACTTIME to allow more time for the server to handle the data.

Tip: Specify the Session option when formatting the packet trace for best results.

Waiting for reply

Timer INACTTIME

Explanation
The FTP client is waiting for an expected reply from the FTP server on the control connection. The timer has expired, or the user has interrupted the FTP client before a reply was received. The reply from the FTP server tells the FTP client whether the previous command was successful or not. When a reply is not received, the FTP client must assume that the command was not successful.

User Response
INACTTIME could be increased to allow the FTP server more time to reply. If the failure recurs, contact the system programmer.

System Programmer Response
For long running jobs, firewalls might time out the control connection due to inactivity. FTPKEEPALIVE can be coded in FTP.DATA to cause the TCP layer to send KeepAlive packets on the control connection. The firewalls can also be configured with longer inactive times. Use a packet trace to determine whether the replies arrive at the FTP client. If the packet trace does not show the FTP reply, determine where the reply is being rejected. Otherwise, contact the IBM Support Center to investigate the packet trace.

Sending command to SOCKS server

Timer INACTTIME

Explanation
The FTP client has timed out sending a command to the SOCKS server. This indicates that the TCP layer is unable to transmit data to the SOCKS server.

User Response
Contact the system programmer.

System Programmer Response
Take a packet trace to investigate the TCP traffic between the two hosts. Use the **locstat** subcommand to determine the IP address of the SOCKS server.

Waiting for reply from SOCKS server

Timer INACTTIME

Explanation
The client is trying to establish a control or data connection to the FTP server through the SOCKS server. The client has sent a connection establishment SOCKS command to the SOCKS server and is waiting for a reply. The FTP client has timed out or been interrupted while waiting for the reply. The SOCKS server might not have replied because it was not processing SOCKS commands in a timely fashion; it was waiting for the remote FTP server to respond, or the SOCKS server did not process the FTP server response in a timely fashion.

User Response

INACTIME can be increased to allow the SOCKS server more time to process commands. If the message occurred while trying to build a data connection through the SOCKS server, issuing the **locsite fwfriendly** subcommand prior to the **put** or **get** subcommand might allow the data connection to be built. If the failure recurs, contact the system programmer.

System Programmer Response

Verify with the administrator of the SOCKS server that the server is receiving the commands and processing them in a timely fashion. The IP address of the SOCKS server can be determined with a **locstat** command. Ensure that the SOCKS server can communicate with the FTP server. Firewalls between the SOCKS server and the FTP server must allow FTP connections and FTP data connections. Take a packet trace to trace the network traffic between the FTP client and SOCKS server. An FTP client trace, enabled by coding DEBUG SOC(2) and DUMP 85 in FTP.DATA, shows the SOCKS commands sent to the server. If a firewall is blocking the data connection, issuing the **locsite fwfriendly** subcommand prior to the **put** or **get** subcommand or specifying FWFRIENDLY TRUE in FTP.DATA might allow the data connection to complete.

Establishing data connection through SOCKS server

Timer MYOPENTIME

Explanation

The FTP client is trying to establish a TCP connection to the SOCKS server so that a data connection can be established to the FTP server. The client has already successfully logged into the FTP server using the SOCKS server. The TCP connection has not completed. The SOCKS server might be too busy to accept new connections in a timely fashion.

User Response

The value of MYOPENTIME can be increased to allow more time for the SOCKS server to accept the connection. If the failure recurs, contact the system programmer.

System Programmer Response

Contact the administrator of the SOCKS server to determine whether the SOCKS server is accepting new connections. Take a packet trace to verify that the SOCKS server is not responding to the connection attempt.

Initial connection to SOCKS server

Timer MYOPENTIME

Explanation

The FTP client is trying to establish a TCP connection to the SOCKS server so that a control connection can be established with the FTP server. The TCP connection has not completed. The SOCKS server might be too busy to accept new connections in a timely fashion.

User Response

Use the **locstat** subcommand to determine the IP address of the SOCKS server. Verify that the SOCKS server is reachable by

pinging the server. Increasing the value of MYOPENTIME allows the SOCKS server more time to accept the connection. If the problem recurs, contact the system programmer.

System Programmer Response
Verify that the SOCKS server is reachable. Contact the administrator of the SOCKS server to determine whether the SOCKS server is accepting new connections. Take a packet trace to determine whether the TCP connection to the SOCKS server completes. Use the **locstat** subcommand to determine the IP address of the SOCKS server; the port number of the SOCKS server is always 1080.

Diagnosing FTP client user exit problems

This section lists and describes possible problems when you install and use the FTP client security user exits.

Message EZA1555I is displayed

The FTP client displays the `EZA1555I CSVDYNEX request failed for user exit exitname, RETURN CODE: x'rc REASON CODE x'rsncode'` message. This occurs when a DES (Dynamic Exits Services) CSVDYNEX request for the FTP client security exits fails. For more information about the return and reason codes for CSVDYNEX macro, see Return and reason codes in z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN.

Message EZA1548I is displayed

The FTP client displays the `EZA1548I User exit exitname module modname abended with abend code x'rc' and abend reason code x'rsncode'` message. This occurs when an FTP client user exit ends abnormally. The MVS console or the system log has the information about the abend code and abend reason code. See z/OS MVS System Codes for an explanation and programmer response.

Message EZA1545I is displayed

The FTP client displays the `EZA1545I Return code rc from user exit exitname module modname is ignored` message in the following situations:

- When the user exit routine is to modify the arguments of commands such as AUTH, EPRT, EPSV, MODE, PBSZ, PORT, PROT, REST, SITE, STRU, TYPE and XLM. Modifying these command arguments leads to unexpected failures because of desynchronization between the FTP client and server, so the modified arguments of these commands are ignored. Correct the user exit routine so that it does not modify the arguments of these commands.
- When the user exit routine is to reject the QUIT command. Correct the exit routine to accept the QUIT command.
- When the user exit routine is to end the FTP client when the QUIT command is from the QUIT subcommand. The FTP client ends as part of QUIT subcommand processing, and it must send QUIT to the server before the client ends. Correct the exit routine to accept the QUIT command. The exit routine does not need to end the FTP client in this circumstance.
- When the user exit routine is to modify the arguments of RETR or STOR during a load module transfer. Changing the arguments of RETR or STOR during a load

module transfer leads to unpredictable results. Correct the exit routine so that it does not modify the arguments of these commands when transferring a load module.

- When the user exit routine is to end the FTP client and the reply is `221 Quit command received`, it is from the QUIT subcommand. The FTP client always ends when it completes QUIT subcommand processing. The exit routine does not need to end the FTP client in this circumstance.
- When the user exit routine is to end the FTP client and the reply starts with reply code in the range 100 - 199. It leads to unpredictable results. Correct the exit routine to accept replies with reply codes in the range 100 - 199.

The FTP client exits unexpectedly, without the user issuing a QUIT subcommand

Message EZA1546I appears

The FTP client displays the `EZA1546I User exit` *exitname* `module` *modname* `ended the FTP client – exit reason code x'`*hexadecimal-rsncode*`' (`*decimal-rsncode*`)` message. Inspect the optional EZA1556I and EZA1546I messages that are displayed by the FTP client to determine why the user exit routine ended the FTP client. No further action is needed if the explanation is satisfactory, and it is appropriate for the exit routine to end the FTP client. Otherwise, correct the user exit routine. See FTP client user exits in z/OS Communications Server: IP Configuration Reference for more information.

Message EZA1549I appears

The FTP client displays the `EZA1549I User exit` *exitname* `module` *modname* `returned an invalid return code` *rc* message. This occurs when the FTP EZAFCCMD client user exit returns a value in register 15 that is not 0, 4, 8 or 12, or when the FTP EZAFCREP client user exit returns a value in register 15 that is not 0 or 12. The systems programmer needs to inspect the messages displayed by the FTP client and correct the user exit routine.

The user exit routine that is installed cannot be called

The FTP client displays the `EZA1555I CSVDYNEX CALL failed for user exit` *exitname*`, RETURN CODE: x'04' REASON CODE: x'00000406'` message. This occurs when no active exit routine is associated with the user exit. This can happen when the exit routine ends abnormally, or is deleted by the operator while an FTP client is logged in to an FTP server. If the exit routine ends abnormally, EZA1548I provides information about the abend code and abend reason code. See "Message EZA1548I is displayed " on page 489 for an explanation and response. Otherwise, see Return and reason codes in z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN for more information.

The FTP subcommand failed

Message EZA1532I appears

The FTP client displays the `EZA1532I User exit` *exitname* `module` *modname* `modified the FTP command arguments` message when the FTP EZAFCCMD client user exit modified the command arguments. If the command arguments are modified to an invalid value, the subcommand might fail. You might need to correct the user exit routine *modname*.

Message EZA1533I appears

The FTP client displays the `EZA1533I User exit` *exitname* `module` *modname* `prevented user` *userid* `from sending command` *command* `– exit reason`

code x'*hexadecimal-rsncode*' (*decimal-rsncode*) message. Inspect the optional EZA1556I and EZA1533I messages that are displayed by the FTP client to determine why the user exit routine rejected the FTP command. If it is appropriate for the exit routine to reject the FTP command, no further action is needed; otherwise, correct the user exit routine *modname*.

Problems starting the client

This topic lists and describes possible problems starting the FTP client.

Enabling or suppressing message EZYFT47I during startup

When the FTP client reads a statement in FTP.DATA that is supported by the z/OS FTP server but not by the FTP client, it issues the message EZYFT47I as a warning. For example, if the client finds an ANONYMOUS statement in FTP.DATA, it issues EZYFT47I for that statement because the ANONYMOUS statement has meaning only for the z/OS FTP server.

If you use the same FTP.DATA configuration file for both client and server, you might want to suppress the EZYFT47I messages. You can prevent the client from issuing this warning by coding a SUPPRESSIGNOREWARNINGS statement in FTP.DATA. Code SUPPRESSIGNOREWARNINGS in FTP.DATA only after you verify all statements in FTP.DATA are correct.

If you require message EZYFT47I for diagnostic purposes, verify no SUPPRESIGNOREWARNINGS statements are coded in FTP.DATA, or else code SUPPRESSIGNOREWARNINGS FALSE in FTP.DATA ahead of those statements you want to debug.

Abends

If the client abends immediately after the FTP command is entered and the following message is displayed, ensure that an OMVS segment is defined for the local TSO user ID or that automatic assignment of unique UNIX identities support is enabled:

```
ftp
 CEE5101C During initialization, the OpenEdition callable service
 BPX1MSS failed. The system return code was 0000000156
     , the reason code was 0B0C00FB . The application will be
     terminated
 IKJ56641I FTP      ENDED DUE TO ERROR+
 READY
```

Incorrect configuration values

Issue the LOCSTAT subcommand to determine the name of the file being used for your local site configuration parameters. If the file you want is not being used, start the FTP client with the **-d** or TRACE options to trace the client as it follows the search order for the FTP.DATA file. For more information about the search order used by the client, see z/OS Communications Server: IP User's Guide and Commands.

Determine whether your FTP.DATA file has sequence numbers. If it does, any statement with an optional parameter omitted picks up the sequence number as the parameter value. For example, the BLKSIZE statement has an optional parameter *size*. If you specify the size, the sequence number is ignored. If you do not specify the size, the system assumes the sequence number is the size, causing an error.

Problems logging in to the server

This topic lists and describes possible problems logging in to the server.

Client ignores SOCKS configuration file

If you suspect that the client consistently ignores the SOCKS configuration file, use the **locstat** subcommand to display the name of the SOCKS configuration file.

- If no SOCKS configuration file name appears in the LOCSTAT output, the client is not configured correctly. Verify that a SOCKSCONFIGFILE statement is in FTP.DATA.
- Inspect the client syslog output for error messages relating to SOCKSCONFIGFILE in FTP.DATA. Use the client DEBUG INT statement to trace client initialization, and look for messages relating to the SOCKS configuration.

The FTP client references the SOCKSCONFIGFILE only when it is connecting to servers with IPv4 IP addresses; it is supposed to ignore the SOCKSCONFIGFILE when logging in to an FTP server with an IPv6 IP address. If you specify the FTP server by DNS name, that name might resolve to an IPv6 address rather than to an IPv4 address. Use the **locstat** subcommand to display the IP address used to log in to the server; the port number of the SOCKS server is always 1080.

Client connects to wrong SOCKS server

If the client connects to a wrong SOCKS server; to a SOCKS server when it should not, or ignores SOCKS configuration file some of the time, use the **locstat** subcommand to display the name of the SOCKS configuration file.

- If the name displayed is not correct, correct the SOCKSCONFIGFILE statement in FTP.DATA.
- If the SOCKS configuration file name displayed by LOCSTAT is correct, inspect the contents of the SOCKS configuration file.

  The client processes the statements in the order they are coded and applies the first statement that specifies the target FTP server. Check and arrange the statements as appropriate, or add a new statement specific to the FTP server at the beginning of the file.

Connection through SOCKS server to FTP server fails

A SOCKS connection involves a connection between the client and SOCKS server, and the SOCKS server and the target server.

When a connection fails, try to isolate the point of failure by checking the following:

- Can client connect to the SOCKS server host?

  Use the client SOC(2) trace and the DUMP 85 trace during connection establishment, and inspect any messages to gain insight into whether the client was able to connect to the SOCKS server.
- Is the link between the client and the SOCKS server good?

  Use ping to test the link.
- Is the SOCKS server active?
- Is the SOCKS server configured to reject the connection?

  Contact the administrator of the SOCKS server for assistance.
- Is the link between the SOCKS server and the FTP server good?

  Ask the administrators of the SOCKS server and the FTP server to verify the link.
- Is the FTP server active and accepting connections?

Contact the administrator of the FTP server. For the z/OS FTP server, activate the trace and check the syslog to determine whether the FTP server received a connection from the SOCKS server on behalf of the client.

Message EZA2589E appears while trying to log in

See "Diagnosing FTP connection and transfer failures with EZA2589E" on page 483.

Server rejects password or password phrase

The z/OS FTP server supports case-sensitive passwords when your RACF administrator has enabled mixed-case passwords. Verify that you have entered the password correctly, and in the correct case. Password phrases used to log in to a z/OS FTP server are always case-sensitive.

If you are using a NETRC file or data set to provide the FTP login password, verify that the password is coded correctly and in the correct case.

If the z/OS FTP server rejects a mixed-case or lower-case password that it formerly accepted, it is possible your RACF administrator has disabled RACF mixed-case password support. In that case, it is not possible to log in with any ID whose password has been set to mixed or lower case. Ask your RACF administrator to reset the password.

If the server rejects a password of more than one token, verify that you enclosed the password phrase with quotation marks when you entered it, or when you coded it in the NETRC file or data set.

FTP clients send passwords to the FTP server as ASCII-encoded data. The server must translate the password or password phrase to its native EBCDIC code page because passwords and password phrases are defined to your z/OS security product as EBCDIC strings. If any character of the password phrase does not translate into the correct character, the password will not match the password assigned to the user ID, and logging in will fail. A character might not translate correctly for the following reasons:

- The client and server are using different ASCII code pages for the control connection. Verify that the client and server are using the same code page for the control connection.
- The client is EBCDIC-based and is using a different EBCDIC code page than the server. Verify that the client and server are using the same native code page.
- The client and server are using matching code pages, but one or more characters in the password phrase does not translate into the correct character. Verify that each character assigned to the password phrase can be translated from EBCDIC to ASCII, and from ASCII back to the original EBCDIC character.

Unknown host error message

The FTP client displays `EZA1551I Unknown Host: <hostname>` if it receives a negative response from the resolver. This occurs when the hostname specified on the FTP command cannot be resolved either by the name server or the local resolution file.

Rule: The FTP client always uses the z/OS UNIX search order for TCPIP.DATA, even when FTP is invoked from TSO.

Use the host IP address instead of the hostname on the FTP command, or see Chapter 39, “Diagnosing resolver problems,” on page 869 for information about diagnosing name server problems.

Problems transferring data

This topic lists and describes possible problems transferring data.

Many data transfer problems that apply to a server apply also to a client. See “Cannot establish conversion between <codeset> and UCS-2” for more information.

Cannot establish conversion between <codeset> and UCS-2

If you invoke the FTP client under TSO, and issue a TYPE U2 or UCS2 subcommand, the following message might be issued:

```
EZA2749E Cannot establish conversion between <codeset>
          and UCS-2.
```

To transfer data encoded in UCS-2 during an FTP session, invoke the FTP command with the _ICONV_UCS2_PREFIX environment variable, specifying the prefix used for your runtime library. Following is an example:

```
FTP ENVAR("_ICONV_UCS2_PREFIX=CEE.OSVIR4") / <host_ip_addr> <port>
```

Secure IPv4 FTP session cannot transfer data through an NAT firewall

If you are using an encrypted FTP control connection, as is the case when using AT-TLS security, and the client sends PASV or PORT to establish a data connection for file transfer, and a NAT (network address translation) firewall exists between the client and server, you might find that while you could sign into the server, you cannot establish the data connection for the transfer. This is because a NAT firewall monitors the FTP control connection as well as the IP headers, changing IP addresses as needed. If the control connection is encrypted, the NAT cannot monitor and change the IP addresses exchanged between the FTP client and server by PASV and PORT.

Use the **locsite** subcommand with the EPSV4 parameter, or code EPSV4 TRUE in FTP.DATA, to direct the client to use EPSV instead of PORT or PASV on IPv4 sessions to establish the data connection. The EPSV command exchanges only port numbers between FTP client and server, so the NAT firewall does not need to translate IP addresses. The server must support EPSV on IPv4 sessions for this solution to be effective.

If the server does not support the EPSV command, you can use the PASSIVEIGNOREADDR configuration option to ignore the IP address that is returned on a PASV command reply and use only the port. For more information about the EPSV command, see RFC 2428. For more information about the LOCSITE subcommand, see z/OS Communications Server: IP User's Guide and Commands and z/OS Communications Server: IP System Administrator's Commands. For more information about the EPSV4 or PASSIVEIGNOREADDR statement in FTP.DATA, see z/OS Communications Server: IP Configuration Reference.

Firewall does not permit FTP client to establish a data connection

You might be able to log in to an FTP server through a firewall, but find you cannot transfer files using a passive data connection. The reason is that the ephemeral ports chosen for the data connection are outside the range of ports permitted by the firewall.

If the client sends EPSV or PASV to the server to start the data connection, FTP is establishing a passive data connection, or is operating in passive mode. In passive mode, the server chooses the ephemeral port for the data connection. Ephemeral port numbers are part of EPSV and PASV replies the server sends to the client. You can configure the z/OS FTP server to use only a specific range of ephemeral ports for the data connection that are compatible with what you have configured for your firewall by coding the PASSIVEDATAPORTS statement in FTP.DATA. See z/OS Communications Server: IP Configuration Reference for information about the PASSIVEDATAPORTS statement. Alternatively, you can restrict the range of ephemeral ports that are available to TCP/IP by coding the EPHEMERALPORTS parameter on the TCPCONFIG configuration statement. See the description of the EPHEMERALPORTS parameter in the TCPCONFIG configuration statement section in z/OS Communications Server: IP Configuration Reference for more information about the EPHEMERALPORTS parameter.

Result: Both FTP and non-FTP applications use ephemeral ports that are within the EPHEMERALPORTS range.

If the client sends PORT or EPRT to the server to start the data connection, the client is said to be establishing an active data connection, or operating in active mode. Active mode FTP is not recommended for sessions through firewalls. Use the **locsite** subcommand with the FWFRIENDLY parameter, or code FWFRIENDLY TRUE in FTP.DATA, to direct the client to operate in passive mode.

Server rejects PORT or EPRT command with 504 replies

Data transfer command sequences that use the PORT or EPRT command fails when the server that receives the PORT or EPRT command is configured to reject all or certain PORT and EPRT commands. The reply code 504 indicates a problem of this nature.

For an ordinary transfer of data between client and server, the z/OS FTP client sends the PORT command to server when:

- The server does not support the EPSV command or the FTP session protocol is IPv4, and
- The client is not configured to be firewall-friendly

You can correct this problem in one of the following ways:

- Make the client firewall-friendly. Do this for the z/OS FTP client by coding FWFRIENDLY TRUE in the client FTP.DATA or by using a LOCSITE FWFRIENDLY subcommand before attempting the data transfer. The client sends EPSV or PASV to the server instead of PORT and the problem is avoided.
- Log in to the server using the server IPv6 address. The client uses EPSV instead of PORT and the problem is avoided.

  Restriction: The server must have an IPv6 address.
- Change the server configuration so that it does not reject PORT or EPRT commands.
- Change the server so that it supports the EPSV command. The z/OS FTP server supports the EPSV command.

To change the client, see z/OS Communications Server: IP User's Guide and Commands for information about the FWFRIENDLY statement and the LOCSITE subcommand.

If you used the **proxy** subcommand to start the transfer, you are transferring data between two servers instead of between client and server. For a transfer of data between two servers, the client must send PORT or EPRT to one of the servers, and PASV or EPSV to the other server. If the server receiving the PORT or EPRT command is configured to reject the PORT or EPRT command, the proxy transfer fails with a 504 reply.

You can fix this problem in one of the following ways:

- Reverse the order in which you open the server connections. That is, if you opened a connection to ServerA and proxy opened a connection to ServerB, open the connection to ServerB and proxy open the connection to ServerA. The client then sends PORT or EPRT to the other server during the proxy transfer. Provided the other server does not also reject PORT or EPRT, this avoids the problem.

  Restriction: If the file you are transferring is a load module, changing the order in which you open server connections does not always cause the client to send PORT or EPRT to the other server.
- Transfer the file to a client, and then to the other server.
- Change the server so that it does not reject PORT and EPRT commands.

The following are z/OS server FTP.DATA statements that can be coded to reject PORT and EPRT commands:

PORTCOMMAND
Reject all PORT and EPRT commands.

PORTCOMMANDPORT
Reject PORT and EPRT commands whose port number argument is a well-known port number.

PORTCOMMANDIPADDR
Reject PORT and EPRT commands whose argument is an IP address that is different from the client's IP address.

See z/OS Communications Server: IP Configuration Guide for more detail.

Message EZA2589E appears when trying to transfer data

See "Diagnosing FTP connection and transfer failures with EZA2589E" on page 483.

Message EZA1771I is displayed

The FTP client displays the `EZA1771I subcommand fails: name is of an unsupported type` message. This message is issued when the data set being transferred is not supported by the MVSGet or MVSPut subcommand. The supported data set types are z/OS physical sequential data set, z/OS partitioned data set, z/OS libraries, and z/OS generation data sets. See the MVSGET subcommand or the MVSput subcommand in z/OS Communications Server: IP User's Guide and Commands for more information.

Message EZA1776I is displayed

The FTP client displays the `EZA1776I subcommand fails: the target generation data set must be a positive relative reference` message. This message is issued when the MVSGet or MVSPut subcommand is being transferred into a generation data set by using a negative or current reference. See z/OS DFSMS Using Data Sets for information about Processing Generation Data Groups.

Message EZA1777I is displayed

The FTP client displays the `EZA1777I subcommand fails: the target generation data set cannot be a library` message. This message is issued when the MVSGet or MVSPut subcommand is transferring a library into a generation data set. But the generation data sets cannot be PDSEs, UNIX files, or VSAM data sets. See z/OS DFSMS Using Data Sets for information about Processing Generation Data Groups.

Message EZA1778I is displayed

The FTP client displays the `EZA1778I the target generation data set must be referenced using its absolute name when the source data set is a PDS` message. When FTP is transferring a PDS into a generation data set, you must reference the generation data set by its absolute data set name. See z/OS DFSMS Using Data Sets for information about Processing Generation Data Groups.

The target data set does not match the source data set using the MVSGet or MVPut subcommand

FTP can determine approximate values for only the primary and secondary allocation and space type, but it uses an allocation that is large enough to contain the data. For complete control over the initial allocation, use the **locsite (site)** subcommand with the **get (put)** subcommand instead of the MVSGet (MVSPut) subcommand.

The MVSGet or MVSPut subcommand fails when transferring a physical sequential extended format data set

When the source data set is of physical sequential extended format, the target data set is allocated as if the DSNTYPE with SYSTEM value was configured. If the system default DSNTYPE value is not EXTREQ or EXTPREF, the source data set might exceed the architecture size limit of the system default DSNTYPE, and the transfer fails.

REStart the failed PDS and library data sets transfer for the MVSGet or MVSPut subcommand

The MVSGet or MVSPut subcommand supports checkpointing for block mode restart of interrupted file transfer for physical sequential data sets only.

Other problems

This topic lists and describes other problems diagnosing FTP connection and transfer failures.

Client PDS member statistics not created or updated

ISPFStats must be set to TRUE to create or update the statistics for the PDS Member when using GET and MGET subcommands. When the PDS directory block is full, PDS member statistics are not updated. Use the **locstat** subcommand to verify that the client's ISPFSstats setting is TRUE. Use the LOCSITE ISPFStats subcommand change the ISPFSstats value. See z/OS Communications Server: IP User's Guide and Commands for information about using the LOCSITE subcommand.

Diagnosing FTP client problems with tracing

You can activate tracing on startup by doing the following:

- Coding DEBUG statements in FTP.DATA. See the DEBUG statement in z/OS Communications Server: IP Configuration Reference for more information.
- Starting the FTP client with the -d command-line option. See z/OS Communications Server: IP User's Guide and Commands for more information about the FTP environment.

Alternatively, you can activate tracing by toggling tracing on or off during an FTP session with the DEBUG subcommand.

You can activate FTP client extended trace at startup by coding one or more DUMP statements in FTP.DATA. See the DUMP statement in z/OS Communications Server: IP Configuration Reference for more information. Alternatively, you can toggle extended tracing on or off during an FTP session with the DUMP subcommand.

The DEBUG and DUMP subcommands activate the general and the extended levels of tracing. The general tracing shows key events in the processing of a subcommand (for example, the opening of a file) and the extended trace shows data areas that are used during processing. The extended trace produces large amounts of output and should be used at the direction of IBM service team. The format of DEBUG allows multiple parameters to be specified on one subcommand. See z/OS Communications Server: IP User's Guide and Commands for the syntax and parameters for the DEBUG and DUMP subcommands.

For example, the following sequence of subcommands would set traces:

```
DEBUG ACC SQL      *Activates the ACC and SQL traces
DEBUG BAS          *Activates the default traces
                   *CMD, INT, FSC, and SOC in addition
                   *to the two already set
DEBUG              *Resets all tracing
```

When running FTP interactively or from a REXX exec, all tracing goes to the terminal unless output is redirected. When running FTP from a TSO batch job, all tracing goes to SYSOUT.

Use the following checklist to diagnosis FTP client problems with tracing:

- Ensure that the user has properly allocated the DDNAME being referred to. The TSO LISTALC STAT HIST command can be helpful in debugging allocations. Also, ensure that the allocations are correct. For example, if a file already exists, the disposition should not be new.
- Ensure that DDNAMEs are only used to refer to local files. For example, `get //DD:FTP01 FILEONE` is not valid because it attempts to use a DDNAME to refer to a host file. If you try to use a DDNAME for a remote file name, the name is sent to the remote host for processing as it is. If the remote host actually has a file named `//DD:FTP01`, then that file would be referred to, but most likely the remote host would reject it as a file name that is not valid.
- To find attempts to access files by DDNAME, look for DD: in FTP trace output as shown below:

```
MF0573 seq_open_file: OSTN -> w,recfm=*,NOSEEK for dd:FTP02
MF0663 seq_open_fle: ddname FTP02 has filename USER1.CCPYXLMT
MF0669 seq_open_file: set DDNAME characteristics- recfm=90, lrecl=128, blksize=614
```

Tip: By using DDNAME support, the user is assuming responsibility for correctly allocating and deallocating the DDNAMEs being used.

Where to find the FTP client trace

The destination of the z/OS FTP client trace depends on the environment in which the client executes as described as follows:

- When the FTP client is invoked interactively from TSO or a REXX exec with an allocated OUTPUT DD, the trace is written to the destination associated with the OUTPUT DD.

- When the FTP client is invoked interactively from a TSO session with no allocated OUTPUT DD, the trace is written to the user's console.
- When the FTP client is invoked interactively from OMVS, the trace is written to the user's console, or it can be written to a file by using the OMVS redirect operand (>).
- When the FTP client is invoked interactively from a REXX exec with no allocated OUTPUT DD, the trace is written to the destination for STDOUT (which might be the user's console).
- When the FTP client is invoked from any application using the FTP Callable Application Programming Interface (API), the trace output is stored in the interface buffer until the application issues a request to retrieve the output. See z/OS Communications Server: IP Programmer's Guide and Reference for a complete description of the FTP Callable API.

Rules: When the FTP client is invoked from a batch job, the following rules apply:

- If the client is invoked directly (EXEC PGM=FTP), the trace is written to the destination associated with the OUTPUT DD.
- If the client is invoked from TSO in batch (EXEC PGM=IKJEFT01), the trace is written to the destination associated with the OUTPUT DD if one exists. Otherwise, the trace is written to the destination associated with the SYSTSPRT DD.
- If the client is invoked from a REXX exec in batch, whether under batch TSO, the trace is written to the destination for the OUTPUT DD (if one exists). Otherwise, the trace is written to the destination for STDOUT (under batch TSO, this might be the SYSTSPRT DD).

Documenting FTP client problems

If the problem is not caused by any of the common errors described in this section, collect the following documentation before calling the IBM Software Support Center. Documentation is divided into the following categories: essential and helpful (but not essential).

- Essential
  - Precise problem description, including client console, expected results, and actual results
  - Include trace in the output by one of these methods:
    - Coding DEBUG statements in the client's FTP.DATA. See z/OS Communications Server: IP Configuration Reference for information about the DEBUG statement.
    - Invoking the FTP client with the -d or TRACE invocation option. See z/OS Communications Server: IP User's Guide and Commands for information about entering the FTP environment.
    - Specifying a DEBUG subcommand in the client command input stream before the affected transfer. Use this option only if the problem does not involve the initial establishment of the FTP control session.
  - FTP.DATA file used by the client.
    - You can use DEBUG ALL to capture all details possible.
    - When activating the trace, use the DEBUG option TIMESTAMPS to time stamp the client trace output. See z/OS Communications Server: IP User's Guide and Commands for information about the DEBUG subcommand and z/OS Communications Server: IP Configuration Reference for information about the DEBUG statement.

- If executing the client in batch, collect all the JES output.

- Helpful:
  - Output from the client **locstat** subcommand
  - Output from the client **stat** subcommand
  - FTP.DATA data set
  - TCPIP.DATA data set
  - If appropriate, sample data to re-create the problem
  - If the FTP.DATA parameter LOGCLIENTERR is TRUE, report the contents of message EZZ9830I. The message is written to the system log and the job log when the client is running in batch and to the user's terminal during an interactive client session.

Chapter 15. Diagnosing z/OS UNIX Telnet daemon (otelnetd) problems

This topic provides diagnostic information for z/OS UNIX Telnet daemon (otelnetd) and contains the following subtopics:

- "Common problems"
- "Debug traces" on page 502

Common problems

The following list describes common problems that you might encounter during execution of the Telnet daemon (otelnetd):

- Diagnostic messages are not being printed to the appropriate file.
  - The diagnostic messages are printed out with the use of syslogd. Ensure that the syslogd is currently active by checking for /etc/syslog.pid.
  - If syslogd is active, ensure that the file where the output is sent is currently allocated. Syslogd creates the file if it is started with the -c runtime option. z/OS UNIX Telnet uses local1.debug for logging messages. Ensure that the syslog.conf file contains an entry for local1.debug or the *.* default file. See the z/OS Communications Server: IP Configuration Guide for more detailed information about syslogd.
  - Ensure also that the specified file exists. Ensure that the permissions on the file are at a minimum 666.
  - Make sure you specify `-t` or `-D all`, or `-t` and `-D all`, as the z/OS UNIX Telnet options in /etc/inetd.conf.
- Use of the arrow keys.

  The arrow keys are not functional in raw mode. This is AIX-like behavior, except that, in AIX, the arrow key produces peculiar characters such as ¬-B on the screen to let the user know not to use arrows. Under rlogin, the cursor moves to where you would want it and correction is allowed, but the shell also treats these characters as part of the original command.
- The keyboard appears to be locked and the user cannot issue commands.

  When executing UNIX-type clients (for example, AIX), if the -k option is specified for Telnet in inetd.conf, Telnet does not allow kludge linemode (see "Setting up the inetd configuration file" on page 599). UNIX clients require character-at-a-time mode to process correctly. If you remove the -k option from the parameters, then the software processes correctly.

  If this does not work, run tracing `-t D all`. Look for `Ept` to determine what the exception conditions are for the `pty`. The number of bytes should equal four. Verify that the exception conditions identified are processed by the TN3270E Telnet server. (Check EZYTE67I messages for more information; see Figure 60 on page 503.)
- EDC5157I An internal error has occurred, rsn=0b8802AF.

  The 2AF of the reason code signifies that the user did not have the proper authority to execute the command. This might result in either the user system having BPX.DAEMON authority set up in its environment, and the proper authorities have not been issued to the user. Another result might be that the user does not have super user authority, which might be required to issue some of these commands.

© Copyright IBM Corp. 2000, 2015

Debug traces

Table 31 describes options that relate to user-controlled trace information.

Table 31. Debug trace options

Option	Sub-Option	Description
-t		Internal tracing, intended to replace the DIAGNOSTICS compile option currently in place within the BSD code.
-D	authentication	Turns on authentication debugging code.
-D	encryption	Turns on encryption debugging code.
-D	options	Prints information about the negotiation of TELNET options.
-D	report	Prints the options information, plus some additional information about what processing is going on.
-D	netdata	Displays the data stream received by telnetd.
-D	ptydata	Displays the data stream written to the pty.
-D	all	Supports all options/report/ptydata/netdata/authentication/encryption options.

Debug trace flows (netdata and ptydata)

When issuing any of the following three trace commands within /etc/inetd.conf (-D ptydata, -D netdata, or -D all), you have the contents in both hexadecimal and ASCII, and the data being sent over the sockets or between the ttys in your syslogd file. If the user is having problems between the parent and the client, try the -D netdata option. If it is between the parent and the child, try the -D ptydata option. If both or either might apply, try the -D all option.

Each set of hexadecimal data is preceded by a three-letter tag. This tag represents the direction the data is flowing from. Figure 59 is a pictorial representation of this flow.

- Int—client to parent
- Ont—parent to client
- Ipt—child to parent
- Opt—parent to child

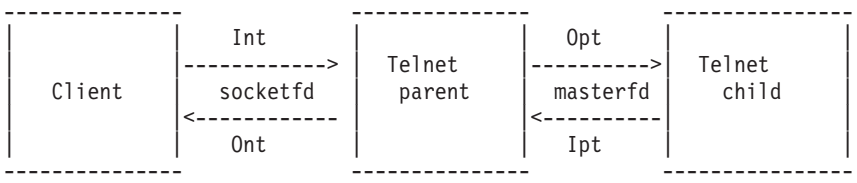

The user types a command on the command line. It flows Int -> Opt. The child responds and the flow is Ipt -> Ont.
Figure 59. Trace between the Telnet client, parent, and child

Debug trace examples (-t -D all)

Figure 60 on page 503 gives an example of the trace generated from **-t -D all**, generated from an AIX Telnet client. A trace explanation follows the figure.

```
 1  EZYTE29I Starting new telnet session.  catfd = 168443936
EZYTO05I Initial EBCDIC codepage = IBM-1047, ascii codepage = ISO8859-1
 2  EZYTE05I Trace 1 Debug 3d keepalive 1 kludgelinemode 0
   hostinfo 1 Registered host 0 linemode 0 multi_proc 0
telnetd: doit(Second_pass=0)
 3  EZYTE11I doit: host_name  laph.raleigh.ibm.com
 4  EZYTE11I doit: IP address 9.37.83.93
EZYTE11I doit: PORT       2504
EZYTE11I doit: host       MVSJ
>>>TELNETD: I support auth type 2 6
>>>TELNETD: I support auth type 2 2
>>>TELNETD: I support auth type 2 0
>>>TELNETD: I will support DES_CFB64
>>>TELNETD: I will support DES_OFB64
telnetd: getterminaltype() auth_level=0
state: send_do(option=37, init=1)
 5  EZYTS04I STATE:send_do: send DO   AUTHENTICATION
 6  EZYTU14I UTILITY: netwrite 3  chars.
 7  EZYTU21I Ont: fffd25 ...
 8  EZYTU03I UTILITY:ttloop read 33 chars.
 9  EZYTU20I Int: fffb25fffd26fffb26fffd03fffb18fffb1ffffb .....................
EZYTU20I Int: 20fffb21fffb22fffb27fffd05 ..............
telrcv() encrypt_output=0
telrcv() decrypt_input =0
 10  EZYTS05I STATE:willoption: receive WILL   AUTHENTICATION
>>>TELNETD: Sending type 2 6
>>>TELNETD: Sending type 2 2
>>>TELNETD: Sending type 2 0
utility: printsub(length=10)
 11  EZYTU17I UTILITY: send suboption
 AUTHENTICATION
 SEND
KERBEROS_V5
CLIENT|MUTUAL|ENCRYPT
KERBEROS_V5
CLIENT|MUTUAL
KERBEROS_V5
CLIENT|ONE-WAY
 12  EZYTS10I STATE:dooption: receive DO   ENCRYPT
 13  EZYTS09I STATE:send_will: send WILL   ENCRYPT
 14  EZYTS05I STATE:willoption: receive WILL   ENCRYPT
state: send_do(option=38, init=0)
 15  EZYTS04I STATE:send_do: send DO   ENCRYPT
utility: printsub(length=6)
 16  EZYTU17I UTILITY: send suboption
 ENCRYPT
 SUPPORT
DES_CFB64
DES_OFB64
 17  EZYTS10I STATE:dooption: receive DO   SUPPRESS GO AHEAD
EZYTS09I STATE:send_will: send WILL   SUPPRESS GO AHEAD
 18  EZYTS05I STATE:willoption: receive WILL   TERMINAL TYPE
state: send_do(option=24, init=0)
 19  EZYTS04I STATE:send_do: send DO   TERMINAL TYPE
EZYTS05I STATE:willoption: receive WILL   NAWS
state: send_do(option=31, init=0)
EZYTS04I STATE:send_do: send DO   NAWS
EZYTS05I STATE:willoption: receive WILL   TSPEED
state: send_do(option=32, init=0)
EZYTS04I STATE:send_do: send DO   TSPEED
EZYTS05I STATE:willoption: receive WILL   LFLOW
state: send_do(option=33, init=0)
```

Figure 60. z/OS UNIX Telnet trace using -t -D all

```
EZYTS04I STATE:send_do: send DO   LFLOW
EZYTS05I STATE:willoption: receive WILL   LINEMODE
state: send_do(option=34, init=0)
EZYTS04I STATE:send_do: send DO   LINEMODE
EZYTS05I STATE:willoption: receive WILL   NEW-ENVIRON
state: send_do(option=39, init=0)
EZYTS04I STATE:send_do: send DO   NEW-ENVIRON
EZYTS10I STATE:dooption: receive DO   STATUS
EZYTS09I STATE:send_will: send WILL   STATUS
>>>TELNETD: in auth_wait.
EZYTU14I UTILITY: netwrite 50  chars.
EZYTU21I Ont: fffa2501020602020200fff0fffb26fffd26fffa ..........0........
EZYTU21I Ont: 26010102fff0fffb03fffd18fffd1ffffd20fffd .....0..............
EZYTU21I Ont: 21fffd22fffd27fffb05 ..........
EZYTU03I UTILITY:ttloop read 512 chars.
EZYTU20I Int: fffa2503757365723532fff0fffa25000206006e ..........0.......>
EZYTU20I Int: 8201c6308201c2a003020105a10302010ea20703 b.F.b.B.....~....s..
EZYTU20I Int: 050020000000a38201126182010e3082010aa003 ......tb../b...b....
EZYTU20I Int: 020105a1101b0e4b52423339302e49424d2e434f ...~...........(..|
EZYTU20I Int: 4da22b3029a003020103a12230201b04686f7374 (s........~......?..
EZYTU20I Int: 1b186d76736a2e7463702e72616c656967682e69 .._........./%......
EZYTU20I Int: 626d2e636f6da381c33081c0a003020101a10302 ._..?_taC.a{.....~..
EZYTU20I Int: 0101a281b30481b01cbcb5a95fd2aa72297fae13 ..sa..a....z^K..."..
EZYTU20I Int: d12bd57b08d13133a485c8a4473c585733ded76e J.N#.J..ueHu......P>
EZYTU20I Int: 711511dfae0a732e0f62329f2c1ec3bd19b35b53 ..............C]..$.
EZYTU20I Int: 58e6ede82efb0c80d525a79d26708f5f78109a85 .W.Y....N.x....^...e
EZYTU20I Int: 54e17ca09ca7a245549229aeecd1a01125338a28 ..@..xs..k...J......
EZYTU20I Int: bb58c3dd526136471c4c0f0688317ac3fefc7f83 ..C../...<..h.:C.."c
EZYTU20I Int: b808ea2bfb64f6ebb0b041cf5edd2f5e43a17f52 ......6.....;..;.~".
EZYTU20I Int: 13a299b7925d3e923df5ef18d690c523e1c35834 .sr.k).k.5..O.E..C..
EZYTU20I Int: 62213fdb0d206b894adec1e1437d9e696d6de8b3 ......,i..A..'..__Y.
EZYTU20I Int: 724c0ed1a48196308193a003020101a2818b0481 .<.Juao.al.....sa..a
EZYTU20I Int: 88f7048e7c7e2b092c0c5301f15d8ed82b92a60d h7..@=......1).Q.kw.
EZYTU20I Int: 9c0524bb740e761ad609ffff09c2a13cbcd952ef ........O....B~..R..
EZYTU20I Int: 704a5a9426a6e2607cfe0d1a3fa9969ba8d20836 ..!m.wS-@....zo.yK..
EZYTU20I Int: b8fdf73528d73abebdb7bbd7135d08e815896c62 ..7..P..]..P.).Y.i%.
EZYTU20I Int: c06e44ce1df73816969e95b77ab5b8d95d2618b9 {>...7..o.n.:..R)...
EZYTU20I Int: 7b2abbe6c6d9adab0320a73aebe5e14f9373503d #..WFR[...x..V.|l.&.
EZYTU20I Int: 18a97c34a5698c5dc56364d871939fee0193fff0 .z@.v..)E..Q.l...l.0
EZYTU20I Int: fffa2605fff0fffa260102fff0fffa1f00780032 .....0......0.......
EZYTU20I Int: fff0fffa2203010300036203 .0..........
telrcv() encrypt_output=0
telrcv() decrypt_input =0
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=10)
 20  EZYTU17I UTILITY: receive suboption
 AUTHENTICATION
 NAME
u
s
e
r
5
2
>>>TELNETD: Got NAME [user52]
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=465)
 21  EZYTU17I UTILITY: receive suboption
 AUTHENTICATION
 IS
KERBEROS_V5
CLIENT|MUTUAL|ENCRYPT
 AUTH 110 130 1 198 48 130 1 194 160 3 2 1 5 161 3 2 1 14 162 7 3 5 0 32 0 0 0
163 130 1 18 97 130 1 14 48 130 1 10 160 3 2 1 5 161 16 27 14 75 82 66 51 57 48
46 73 66 77 46 67 79 77 162 43 48 41 160 3 2 1 3 161 34 48 32 27 4 104 111 115
116 27 24 109 118 115 106 46 116 99 112 46 114 97 108 101 105 103 104 46 105 98
 109 46 99 111 109 163 129 195 48 129 192 160 3 2 1 1 161 3 2 1 1 162 129 179 4
```

```
 129 176 28 188 181 169 95 210 170 114 41 127 174 19 209 43 213 123 8 209 49 51
 164 133 200 164 71 60 88 87 51 222
>>>REPLY:2: [3] (91)
 6f 59 30 57 a0 03 02 01 05 a1 03 02 01 0f a2 4b
>>>REPLY:2: [2] (21)
 75 73 65 72 35 32 40 4b 52 42 33 39 30 2e 49 42
 22  telnetd: Kerberos5 identifies him as ``user52@KRB390.IBM.COM''
EZYTU14I UTILITY: netwrite 130  chars.
EZYTU21I Ont: fffa25020206036f593057a003020105a1030201 .......?........~...
EZYTU21I Ont: 0fa24b3049a003020101a2420440117ac36284cd .s........s.. .:C.d.
EZYTU21I Ont: 65384025a9ee70511777fd91aa4c367edd20162f .. .z......j.<.=....
EZYTU21I Ont: 736d40b6e61fae60d2c74c25aa610dcd10526eea ._ .W..-KG<../....>.
EZYTU21I Ont: b096d7e7ed90a06f7a595cddbe92369be741fff0 .oPX...?:.*..k..X..0
EZYTU21I Ont: fffa250202060275736572353240 4b5242333930 ..... ......
EZYTU21I Ont: 2e49424d2e434f4dfff0 ...(..|(.0
utility: printsub(length=4)
 23  EZYTU17I UTILITY: receive suboption
 ENCRYPT
 REQUEST-START
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=5)
 24  EZYTU17I UTILITY: receive suboption
 ENCRYPT
 SUPPORT
DES_OFB64
>>>TELNETD: He is supporting DES_OFB64 (2)
Creating new feed
utility: printsub(length=14)
 25  EZYTU17I UTILITY: send suboption
 ENCRYPT
 IS
DES_OFB64
OFB64_IV 123 117 204 223 5 21 98 2
>>>TELNETD: (*ep->start)() returned 6
EZYTS17I Defer suboption negotiation
EZYTU14I UTILITY: netwrite 16  chars.
EZYTU21I Ont: fffa260002017b75ccdf05156202fff0 ......#........0
utility: printsub(length=7)
EZYTU17I UTILITY: receive suboption
 NAWS
  0 120 (120)
  0 50 (50)
auth_wait: auth_context a080a30, validuser 3
auth_wait: auth_level 0
telnetd: authtelnet client name: user52 auth name: user52
state: send_do(option=38, init=0)
EZYTS04I STATE:send_do: send DO   ENCRYPT
state: send_do(option=24, init=1)
state: send_do(option=32, init=1)
state: send_do(option=35, init=1)
EZYTS04I STATE:send_do: send DO   XDISPLOC
state: send_do(option=39, init=1)
state: send_do(option=36, init=1)
EZYTS04I STATE:send_do: send DO   OLD-ENVIRON
EZYTS09I STATE:send_will: send WILL   ECHO
EZYTU14I UTILITY: netwrite 12  chars.
EZYTU21I Ont: fffd26fffd23fffd24fffb01 ............
EZYTU03I UTILITY:ttloop read 47 chars.
EZYTU20I Int: 04020f05030007621c08020409421a0a02080b02 .....................
EZYTU20I Int: 150c02170d02120e02160f021110021311020012 .....................
EZYTU20I Int: 0200fff0fffd03 ...0...
telrcv() encrypt_output=0
telrcv() decrypt_input =0
EZYTS17I Defer suboption negotiation
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=52)
EZYTU17I UTILITY: receive suboption
```

```
 LINEMODE
 SLC
 SYNCH
 DEFAULT
 0;
 IP
 VARIABLE
|FLUSHIN|FLUSHOUT
 3;
 AO
 VARIABLE
 15;
 AYT
 DEFAULT
 0;
 ABORT
 VARIABLE
|FLUSHIN|FLUSHOUT
 28;
 EOF
 VARIABLE
 4;
 SUSP
 VARIABLE
|FLUSHIN
 26;
 EC
 VARIABLE
 8;
 EL
 VARIABLE
 21;
 EW
 VARIABLE
 23;
 RP
 VARIABLE
 18;
 LNEXT
 VARIABLE
 22;
 XON
VARIABLE
 17;
 XOFF
 VARIABLE
 19;
 FORW1
 VARIABLE
 0;
 FORW2
 VARIABLE
 0;
EZYTS10I STATE:dooption: receive DO   SUPPRESS GO AHEAD
EZYTU03I UTILITY:ttloop read 16 chars.
EZYTU20I Int: fffa26000201094321d752693162fff0 .........P.....0
telrcv() encrypt_output=0
telrcv() decrypt_input =0
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=14)
 26  EZYTU17I UTILITY: receive suboption
 ENCRYPT
 IS
DES_OFB64
OFB64_IV 9 67 33 215 82 105 49 98
CFB64: initial vector received
Initializing Decrypt stream
```

```
utility: printsub(length=6)
 27  EZYTU17I UTILITY: send suboption
 ENCRYPT
 REPLY
DES_OFB64
OFB64_IV_OK
(*ep->is)(a09abc3, 9) returned MORE_TO_DO (7)
EZYTU14I UTILITY: netwrite 8  chars.
EZYTU21I Ont: fffa26020202fff0 .......0
EZYTU03I UTILITY:ttloop read 17 chars.
EZYTU20I Int: fffa26020202fff0fffc23fffc24fffd01 .......0.........
telrcv() encrypt_output=0
telrcv() decrypt_input =0
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=6)
 28  EZYTU17I UTILITY: receive suboption
 ENCRYPT
 REPLY
DES_OFB64
OFB64_IV_OK
utility: printsub(length=5)
 29  EZYTU17I UTILITY: send suboption
 ENCRYPT
 ENC_KEYID
 0
(*ep->reply)(a09abc3, 1) returned MORE_TO_DO (4)
>>>TELNETD: encrypt_reply returned 4
 30  EZYTS08I STATE:wontoption: receive WON'T   XDISPLOC
EZYTS08I STATE:wontoption: receive WON'T   OLD-ENVIRON
EZYTS10I STATE:dooption: receive DO   ECHO
>>>TELNETD: in encrypt_wait
EZYTU14I UTILITY: netwrite 7  chars.
EZYTU21I Ont: fffa260700fff0 ......0
EZYTU03I UTILITY:ttloop read 7 chars.
EZYTU20I Int: fffa260700fff0 ......0
telrcv() encrypt_output=0
telrcv() decrypt_input =0
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=5)
 31  EZYTU17I UTILITY: receive suboption
 ENCRYPT
 ENC_KEYID
 0
utility: printsub(length=5)
 29  EZYTU17I UTILITY: send suboption
 ENCRYPT
 DEC_KEYID
 0
EZYTU14I UTILITY: netwrite 7  chars.
EZYTU21I Ont: fffa260800fff0 ......0
EZYTU03I UTILITY:ttloop read 14 chars.
EZYTU20I Int: fffa260800fff0fffa260300fff0 ......0......0
telrcv() encrypt_output=0
telrcv() decrypt_input =0
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=5)
 31  EZYTU17I UTILITY: receive suboption
 ENCRYPT
 DEC_KEYID
 0
>>>TELNETD: Encrypt start: initial negotiation in progress (0) DES_OFB64
utility: printsub(length=5)
 32  EZYTU17I UTILITY: send suboption
 ENCRYPT
 START
>>>TELNETD: Started to encrypt output with type DES_OFB64
EZYTU14I UTILITY: netwrite 7  chars.
```

```
EZYTU21I Ont: fffa260300fff0 ......0
utility: printsub(length=5)
 33  EZYTU17I UTILITY: receive suboption
 ENCRYPT
 START
>>>TELNETD: Start to decrypt input with type DES_OFB64
utility: printsub(length=5)
EZYTU17I UTILITY: send suboption
 ENCRYPT
 REQUEST-START
>>>TELNETD: Request input to be encrypted
>>>TELNETD: Encrypt start: initial negotiation in progress (0) DES_OFB64
utility: printsub(length=5)
EZYTU17I UTILITY: send suboption
 ENCRYPT
 START
>>>TELNETD: Started to encrypt output with type DES_OFB64
telnetd: getterminaltype() auth_negotiated=1
utility: printsub(length=4)
 34  EZYTU17I UTILITY: send suboption
 TERMINAL-TYPE
  SEND
EZYTU14I UTILITY: netwrite 32  chars.
EZYTU21I Ont: b306079a3675d15d45511c7172c0579b93f4b0ac ......J).....{..l4..
EZYTU21I Ont: afdc1e09ae1fe760ed4d59b4 ......X-.(..
EZYTU03I UTILITY:ttloop read 51 chars.
EZYTU20I Int: b306079c3675d15d45571f48bb00984d8ac37f6c ......J)......q(.C"%
EZYTU20I Int: afd6c6f276ef18cfa609f445b8b52a92fb4c2a25 .OF2....w.4....k.<..
EZYTU20I Int: ca539aa8f2401960713bc5 ...y2 .-..E
telrcv() encrypt_output=0xA00C848
telrcv() decrypt_input =0xA00C6C0
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=5)
EZYTU17I UTILITY: receive suboption
 ENCRYPT
 START
>>>TELNETD: Start to decrypt input with type DES_OFB64
EZYTS17I Defer suboption negotiation
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=13)
EZYTU17I UTILITY: receive suboption
 TERMINAL-SPEED
  IS 9600,9600
EZYTS17I Defer suboption negotiation
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=16)
EZYTU17I UTILITY: receive suboption
 NEW-ENVIRON
 IS
 VAR
U
S
E
R
 VALUE
u
s
e
r
5
2
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=9)
EZYTU17I UTILITY: receive suboption
 TERMINAL-TYPE
  IS XTERM
EZYTE10I terminaltypeok: call tgetent (buf, XTERM)
```

```
EZYTE51W terminaltypeok: Tgetent failure  EDC5129I No such file or directory.
rsn = 0594003D
 35  EZYTE10I terminaltypeok: call tgetent (buf, xterm)
telnetd: getterminaltype() return 3
 EZYTE22I herald() entered for /etc/otelnetd.banner
 36   EZYTE88E herald: stat error on /etc/otelnetd.banner
     EDC5129I No such file or directory. rsn = 053B006C
     EZYTO01I Int: 75 .
     EZYTO02I Ont: 49 .
     EZYTO01I Int: 73 .
     EZYTO02I Ont: b7 .
     EZYTO01I Int: 65 .
     EZYTO02I Ont: 50 &
     EZYTO01I Int: 72 .
     EZYTO02I Ont: 77 .
     EZYTO01I Int: 35 .
     EZYTO02I Ont: 91 j
     EZYTO01I Int: 32 .
     EZYTO02I Ont: f6 6
     EZYTO01I Int: 0d .
     EZYTO02I Ont: e1 .
     EZYTE59I read_pw: Character ignored 0
 37   EZYTO04I lusername  = user52
     telnetd: krb name: user52, user: user52
     EZYTE22I herald()
     EZYTE22I herald() entered for /etc/banner
 38   EZYTE88E herald: stat error on /etc/banner
     EDC5129I No such file or directory. rsn = 053B006C
     EZYTE16I uid = 52, gid = 5
     telnsave: mallocTelnetSave() rc=0
     telnetd: doit() subcount=96
     telnetd: doit() execvp()
     EZYTU34I id 30002 pri 3 call catopen(tnmsgs.cat,0) code 81 reason 053B006C
     h_errno N/A
     telnetd: main() -y getsubopt(tSave=2137884232)
     telnsave: freeTelnetSave() rc=0
     EZYTO11I DInt: fffa1f00780032fff0fffa220301030003620304 ........0...........
     EZYTO11I DInt: 020f05030007621c08020409421a0a02080b0215 .....................
     EZYTO11I DInt: 0c02170d02120e02160f02111002131102001202 .....................
     EZYTO11I DInt: 00fff0fffa2000393630302c39363030fff0fffa ..0.............0..
     EZYTO11I DInt: 270000555345520175736572353 2fff0 ..............0
     telnetd: doit(Second_pass=1)
     EZYTY02I GETPTY: open of /dev/ptyp   EDC5114I Resource busy. rsn = 020A0155
     EZYTY02I GETPTY: open of /dev/ptyp   EDC5114I Resource busy. rsn = 020A0155
     EZYTY05I GETPTY: slave fd = 9 , masterfd = 8
     telnetd: doit() deferred_processing=1
 39   EZYTS15I STATE:dooption:deferred receive DO   ECHO
     EZYTO09I options(1) = 3 .
     EZYTS15I STATE:dooption:deferred receive DO   SUPPRESS GO AHEAD
     EZYTO09I options(3) = 3 .
     EZYTS15I STATE:dooption:deferred receive DO   STATUS
     EZYTO09I options(5) = 3 .
 39   EZYTS16I STATE:willoption:deferred receive WILL   TERMINAL TYPE
     EZYTO09I options(24) = 12 .
     EZYTS16I STATE:willoption:deferred receive WILL   NAWS
     EZYTO09I options(31) = 12 .
     EZYTS16I STATE:willoption:deferred receive WILL   TSPEED
     EZYTO09I options(32) = 12 .
     EZYTS16I STATE:willoption:deferred receive WILL   LFLOW
     EZYTO09I options(33) = 12 .
     EZYTS16I STATE:willoption:deferred receive WILL   LINEMODE
     EZYTU14I UTILITY: netwrite 13  chars.
     EZYTU21I Ont: b7c2d637e5d127b4aeced4cbad .BO.VJ....M.[
     EZYTO09I options(34) = 12 .
     EZYTS16I STATE:willoption:deferred receive WILL   AUTHENTICATION
     EZYTO09I options(37) = 12 .
     EZYTS15I STATE:dooption:deferred receive DO   ENCRYPT
```

```
     EZYTO09I options(38) = 15 .
     EZYTS16I STATE:willoption:deferred receive WILL   NEW-ENVIRON
     EZYTO09I options(39) = 12 .
     telrcv() encrypt_output=0xA00C848
     telrcv() decrypt_input =0xA00C6C0
     EZYTS18I Process deferred suboption negotiation
     EZYTU14I UTILITY: netwrite 0  chars.
     utility: printsub(length=7)
     EZYTU17I UTILITY: receive suboption
      NAWS
        0 120 (120)
        0 50 (50)
     EZYTS18I Process deferred suboption negotiation
     EZYTU14I UTILITY: netwrite 0  chars.
     utility: printsub(length=52)
EZYTU17I UTILITY: receive suboption
 LINEMODE
 SLC
 SYNCH
 DEFAULT
 0;
 IP
 VARIABLE
|FLUSHIN|FLUSHOUT
 3;
 AO
 VARIABLE
 15;
 AYT
 DEFAULT
 0;
 ABORT
 VARIABLE
|FLUSHIN|FLUSHOUT
 28;
 EOF
 VARIABLE
 4;
 SUSP
 VARIABLE
|FLUSHIN
 26;
 EC
 VARIABLE
 8;
 EL
 VARIABLE
 21;
 EW
 VARIABLE
 23;
 RP
 VARIABLE
 18;
 LNEXT
 VARIABLE
 22;
 XON
 VARIABLE
 17;
 XOFF
 VARIABLE
 19;
 FORW1
 VARIABLE
 0;
 FORW2
```

```
 VARIABLE
 0;
EZYTS18I Process deferred suboption negotiation
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=13)
EZYTU17I UTILITY: receive suboption
 TERMINAL-SPEED
  IS 9600,9600
EZYTS18I Process deferred suboption negotiation
EZYTU14I UTILITY: netwrite 0  chars.
utility: printsub(length=16)
EZYTU17I UTILITY: receive suboption
 NEW-ENVIRON
 IS
 VAR
U
S
E
R
 VALUE
u
s
e
r
5
2
telnetd: doit() deferred_processing=0
state: send_do(option=31, init=1)
state: send_do(option=33, init=1)
telrcv() encrypt_output=0xA00C848
telrcv() decrypt_input =0xA00C6C0
state: send_do(option=0, init=1)
EZYTS04I STATE:send_do: send DO   BINARY
EZYTS07I STATE:send_dont: send DON'T   LINEMODE
EZYTU14I UTILITY: netwrite 66  chars.
EZYTU21I Ont: 2e925cb2dbf210c689e053fbad994f522421dbb7 .k*..2.Fi\..[r|.....
EZYTU21I Ont: 062ef6d290bf59f7e40600bc4c43f4eef139b405 ..6K...7U...<.4.1...
EZYTU21I Ont: a59b84b3fc185a609644499e56c69fe7790b6e7e v.d...!-o....F.X`.>=
EZYTU21I Ont: 8cdd2ede8d42 ......
utility: printsub(length=52)
EZYTU17I UTILITY: send suboption
 LINEMODE
 SLC
 SYNCH
 NOSUPPORT
 0;
 IP
 VARIABLE
|ACK|FLUSHIN|FLUSHOUT
 3;
 AO
 VARIABLE
|ACK
 15;
 AYT
 NOSUPPORT
 0;
 ABORT
 VARIABLE
|ACK|FLUSHIN|FLUSHOUT
 28;
 EOF
 VARIABLE
|ACK
 4;
 SUSP
 VARIABLE
```

```
|ACK|FLUSHIN
 26;
 EC
 VARIABLE
|ACK
 8;
 EL
 VARIABLE
|ACK
 21;
 EW
 VARIABLE
|ACK
 23;
 RP
 VARIABLE
|ACK
 18;
 LNEXT
 VARIABLE
|ACK
 22;
 XON
 VARIABLE
|ACK
 17;
 XOFF
 VARIABLE
|ACK
 19;
 FORW1
 VARIABLE
|ACK
 0;
 FORW2
 NOSUPPORT
 0;
EZYTU14I UTILITY: netwrite 0  chars.
EZYTE66I PROTOCOL:lmodetype=4, linemode=0, uselinemode=0
 40  EZYTY08I argv_fsum(0) =  fomtlinp
EZYTY08I argv_fsum(1) =  *4OurhrEa)R0,H/h
EZYTY08I argv_fsum(2) =
EZYTY08I argv_fsum(3) =  0
EZYTY08I argv_fsum(4) =  8
EZYTY08I argv_fsum(5) =  9
EZYTY08I argv_fsum(6) =  0
EZYTY08I argv_fsum(7) =  0
EZYTY08I argv_fsum(8) =  6
EZYTY08I argv_fsum(9) =  80
EZYTY08I argv_fsum(10) = laph.raleigh.ibm.com
EZYTY08I argv_fsum(11) = xterm
EZYTY08I argv_fsum(12) =
EZYTY08I argv_fsum(13) =
EZYTY08I argv_fsum(14) =
EZYTY08I argv_fsum(15) =
EZYTY08I argv_fsum(16) = 1
EZYTY08I inherit flag =  40000000
EZYTY09I login_tty: spawnp fsumoclp 33
 41  EZYTE67I S(nfd):socketfd..ibits=00000001 obits=00000000 ebits=00000000
         S(nfd) pty..ibits=00000000 obits=00000000 ebits=00000100
 42  EZYTE68I Ept: #bytes = 4 pkcontrol(cntl) 1003
EZYTE69I PROTOCOL: cntl = 1003
EZYTE65I PROTOCOL: send IAC Data Mark.  DMARK
```

Following are short descriptions of the numbered items in the trace:

1 EZYTE29I indicates the start of a new z/OS UNIX Telnet client session.

2 EZYTE05I indicates what options were specified in /etc/inetd.conf for z/OS UNIX Telnet.

3 EZYTE11I indicates the resolved host name (from the client).

4 EZYTE11I shows the IP address of the z/OS UNIX Telnet client.

5 EZYTS04I indicates otelnetd agrees to send and receive authentication information.

6 EZYTU14I traces netwrites (writes to the client terminal).

7 EZYTU21I traces data from parent to client; that is, z/OS UNIX Telnet to the client terminal.

8 EZYTU03I indicates the number of bytes read from the client by z/OS UNIX Telnet.

9 EZYTU20I traces data from the client to the parent (z/OS UNIX Telnet server).

10 EZYTS05I indicates the client agrees to send and receive authentication information.

11 EZYTU17I shows otelnetd requesting that the client send authentication information for Kerberos Version 5.

12 EZYTS10I indicates the client agrees to receive encrypted data.

13 EZYTS09I indicates otelnetd agrees to send encrypted data.

14 EZYTS05I indicates the client agrees to send encrypted data.

15 EZYTS04I indicates otelnetd agrees to receive encrypted data.

16 EZYTU17I shows which types of encryption otelnetd supports when receiving data.

17 EZYTS10I shows the terminal option negotiation the client has sent/received.

18 EZYTS05I shows the terminal option negotiation the client has sent/received.

19 EZYTS04I indicates the terminal negotiation options sent to the client by the z/OS UNIX Telnet server.

20 EZYTU17I shows the account name on otelnetd that the client wishes to be authorized to use.

21 EZYTU17I shows the client authentication information for Kerberos Version 5.

22 Shows the Kerberos Version 5 principal of the user logging in.

23 EZYTU17I shows the client requesting that otelnetd enable encryption as soon as the initialization is completed.

24 EZYTU17I shows which types of encryption the client supports when receiving data.

25 EZYTU17I shows otelnetd sending to the client the type of encryption to use for the data stream (otelnetd to client) and the initial encryption data.

26 EZYTU17I shows otelnetd receiving from the client the type of encryption to use for the data stream (client to otelnetd) and the initial encryption data.

27 EZYTU17I shows otelnetd acknowledging receipt of the initial encryption data from the client.

28 EZYTU17I shows the client acknowledging receipt of the initial encryption data from otelnetd.

29 EZYTU17I shows otelnetd verifying its keyids.

30 EZYTS08I shows the terminal option negotiation the client has sent/received.

31 EZYTU17I shows the client verifying its keyids.

32 EZYTU17I shows all data following this command in the data stream (otelnetd to client) are encrypted using the previously negotiated method of data encryption.

33 EZYTU17I shows all data following this command in the data stream (client to otelnetd) are encrypted via the previously negotiated method of data encryption.

34 EZYTU17I traces z/OS UNIX Telnet sending terminal negotiation suboptions to the client.

35 EZYTE10I traces the call to tgetent(), which determines client terminal type.

36 EZYTE88E indicates no /etc/otelnetd.banner file was found.

37 EZYTO04I shows the user name with which the telnet client logged in.

38 EZYTE88E indicates no /etc/banner file was found.

39 EZYTS15I and EZYTS16I show that a state change was processed due to options/responses received from the client.

40 EZYTY08I traces the parameters passed to the spawned/forked child address space where the OMVS shell runs.

41 EZYTE67I traces the socket sets to show whether input/ibits, output/obits, or exception/ebits data has been received.

42 EZYTE68I shows exception data received on the parent/child connection.

Cleaning up the utmp entries left from dead processes

Assuming that you have the suggested /etc/rc script, the utmpx file is cleaned up each time the S OMVS command is issued. The utmpx file should not normally need cleaning up, as each terminal slot should be reused the next time someone logs on with that terminal.

Although during normal processing the utmp entries are cleaned up, there are the occasional incidents where zombies are created, or the user might have terminated the session abnormally. When this occurs, the utmp entry for that user remains in the /etc/utmpx file until it is cleared out. There is an associated tty reserved for every entry in the /etc/utmpx file including the zombie entries. For dead entries, these ttys are not available for reuse until someone under superuser erases the /etc/utmpx file.

Tip: If you erase the file while someone is logged on, the next logoff reports not finding the utmpx entry for the user. This can be seen with a waitpid failure during that user cleanup.

Chapter 16. Diagnosing Telnet problems

This topic describes how to diagnosis Telnet problems, and contains the following subtopics:

- "General TN3270E Telnet server information"
- "TN3270E Telnet server definitions"
- "Diagnosing TN3270E Telnet server problems" on page 516
- "General Telnet client information" on page 524
- "Telnet client definitions" on page 524
- "Diagnosing Telnet client problems" on page 524
- "Telnet client traces" on page 528

General TN3270E Telnet server information

The Telnet protocol provides a standardized interface, through which a program on one host (the Telnet client) can access the resources of another host (the TN3270E Telnet server) as though the client were a local terminal connected to the server host.

Telnet protocol is based on the concept of a Network Virtual Terminal (NVT) and the principle of negotiated options.

An NVT is an imaginary device, providing the necessary basic structures for a standard terminal. Each host client represents an imaginary device with certain terminal characteristics that the host server can support.

The principle of negotiated options is used by the Telnet protocol because many clients and hosts require additional services beyond the base services. Various options can be negotiated. Server and client use a set of conventions to establish operational characteristics for their Telnet connection by means of the DO, DON'T, WILL, WON'T mechanism that is discussed in "Telnet commands and options" on page 532.

Component event tracing is done under the SYSTCPIP component. A subset of trace options and a subset of IPCS commands are available to Telnet. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 and Chapter 6, "IPCS subcommands for TCP/IP," on page 173 for details.

TN3270E Telnet server definitions

Telnet LUs must be defined correctly to both VTAM and Telnet. A VTAM APPL definition statement is needed for each Telnet LU that is used. Model application definitions can also be used. See the z/OS Communications Server: SNA Resource Definition Reference for detailed information about these definitions. A corresponding LU must be specified in the BEGINVTAM section of the PROFILE data set. See the z/OS Communications Server: IP Configuration Reference for detailed information about these definitions.

Note: All default 3270 LOGMODE entries from the table of Telnet device name parameters in the z/OS Communications Server: IP Configuration Reference are

© Copyright IBM Corp. 2000, 2015

for non-SNA sessions. You must code device types and the needed LOGMODE entries for SNA sessions. All default 3270E LOGMODES are for SNA sessions.

Diagnosing TN3270E Telnet server problems

Problems with Telnet are generally reported under one of the following categories:

- Abends
- Logon problems
- Session hangs
- Incorrect output
- Session outages

Use the information provided in the following sections for problem determination and diagnosis of errors reported against Telnet.

Abends (server)

An abend during Telnet processing should result in messages and error-related information sent to the MVS system console. A dump of the error is needed unless the symptoms already match a known problem.

Documentation

Code a SYSMDUMP DD or SYSABEND DD statement in the PROC used to start Telnet to ensure that a useful dump is obtained in the event of an abend.

Analysis

See z/OS Problem Management or see Chapter 3, “Diagnosing abends, loops, and hangs,” on page 21, for debugging dumps produced during TCP/IP processing.

Logon problems (server)

Telnet logon problems are reported when clients are unable to connect to the host application. Generally, this type of problem is caused by an error in the configuration or definitions (either in VTAM or Telnet).

If the problem can be re-created, use the DEBUG DETAIL parameter to gather diagnostic messages or trace information. See z/OS Communications Server: IP Configuration Guide for details.

Documentation

The following documentation should be available for initial diagnosis of Telnet login problems:

- Console Log of error messages issued by both Telnet and VTAM
- PROFILE data set
- VTAM APPL definitions for Telnet LUs

More documentation that might be needed is discussed in the following analysis section.

Steps for analyzing logon problems (server)

Determine why clients are unable to connect to a host application.

About this task

Table 32 shows symptoms of login problems and refers to the steps needed for initial diagnosis of the error. The information following the chart and associated information can be used for extended diagnosis, if the problem persists.

Table 32. Telnet login problems

Login problem	Analysis steps
No LUs available	1, 2, 6, 10, 13
OPEN failure	1, 2, 3, 4, 5, 6, 7, 8, 9, 10
x-clock (Telnet solicitor panel)	1, 2, 3, 4, 5, 6, 7, 10
x-clock (blank screen)	1, 2, 3, 6, 7, 8, 10, 12
x-clock (application panel)	7, 8, 10
Incorrect USSMSG or DEFAULTAPPL	3, 4, 5, 6, 1, 13, 14

Procedure

Perform the following steps:

1. Have VTAM APPL definition statements been coded correctly?

   Note: There must be a VTAM definition statement or model application name for each LU coded in the PROFILE data set.
2. Is the VTAM node containing the Telnet LU definitions active?
3. Is there a DEFAULTAPPL coded in the PROFILE data set?
4. Is the host application (or DEFAULTAPPL) active?
5. Is there an ALLOWAPPL statement coded that includes the requested application?
6. Have comment delimiters been added or removed as needed in the BEGINVTAM section of the PROFILE data set?
7. Have correct LOGMODEs (or required overrides for SNA) been coded in the PROFILE data set?
8. Does the host application have BIND (session parameter) requirements that are not met by the specified LOGMODE?
9. Is the MSG07 parameter coded in the PROFILE.TCP data set?

   Note: MSG07 returns information to the end user indicating the reason for the failure.
10. Are any abends (in VTAM, host application, or TCP/IP) indicated on the MVS system console?

    Note: If an abend occurred, see the section on abends to continue investigation of the problem.
11. Check the PROFILE data set for the IP to LU mapping.
12. Is an SSL client attempting to connect to a basic port or is a basic client trying to connect to an SSL port?
13. Use the D TCPIP,,T,PROFILE,DETAIL command to view the active profile definitions.
14. Determine whether USSTCP within the TCPIP PROFILE points at the correct USSTAB, because this could also cause an incorrect USSMSG to be displayed.

Results

If the problem still occurs after following the preceding procedure and making any needed changes, obtain the following documentation:

- TELNET display of the LUNAME or CONN ID of affected client, for example, D TCPIP,,T,CONN,LUN=luname.
- VTAM DISPLAY of Telnet LU.
- VTAM DISPLAY of the target host application.
- Activate DEBUG DETAIL and review additional diagnostic information this function provides.

For information about the Telnet Display command options, see the z/OS Communications Server: IP System Administrator's Commands.

The following documentation might also be needed in some cases, but it is suggested that your IBM Software Support Center be contacted before this documentation is obtained:

- TCP/IP packet trace and CTRACE with TELNET option filtered on the IP address of the failing client.
- VTAM buffer trace of the Telnet LU.
- VTAM INTERNAL TRACE (VIT) with options (API,PIU,MSG,PSS,NRM,SSCP).
- Dump of the Telnet address space. To capture the necessary areas of storage in the DUMP command, include:

  ```
  SDATA=(CSA,LSQA,PSA,RGN,SQA,SUM,SWA,TRT,LPA)
  ```

For information about obtaining VTAM traces, see z/OS Communications Server: SNA Operation or to z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT for your release. Instructions on obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids for your release of MVS.

Session hangs (server)

This topic discusses diagnosis of a hang after a session has been successfully connected. A hang would be indicated by the keyboard remaining locked on the client side of the session, with no data being sent to or received from the server host.

If a problem is recreatable, you can use the DEBUG TRACE parameter. See the z/OS Communications Server: IP Configuration Guide for details.

Documentation

To determine the cause of a Telnet session hang, the following documentation is required:

- CTRACE specifying the TELNET option filtered on the IP address of the failing client.
- In some cases a VTAM buffer trace of the Telnet LU might be needed.
- Information about what was seen at the client screen.

Steps for analyzing session hangs (server)

Determine why a session hangs after it was successfully connected.

About this task

The preceding traces are essential to finding the reason for the session hang. Data entered at the client terminal is sent to Telnet on the TCP/IP connection. The TCP/IP packet trace shows the data arriving at or leaving the stack. CTRACE with the option Telnet specified shows the data coming into and out of Telnet (from both the stack and VTAM). Some processing steps during this time are also included in the trace. The CTRACE with Telnet option shows what Telnet does with this data.

The VTAM buffer trace shows the data as received by VTAM to be forwarded to the host application. Following the data flow through the traces between VTAM, TCP/IP, and Telnet provides an indication of where the problem is occurring.

The following list suggests information to check in the traces. See z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT or *SNA Formats* for more information about VTAM buffer trace output.

Procedure

Perform the following steps:

1. Does the packet trace show data passed to TCP/IP? If not, the problem is in client or emulator code. If data is in the trace, continue with Step 2.
2. Does CTRACE with TELNET option show data passed to Telnet? The TELNET option shows data coming into Telnet from the stack and also going out to VTAM (the reverse for outbound data). If not, the error is in the TCP/IP platform code. Otherwise, continue with Step 3.
3. Does VTAM buffer trace show data passed from Telnet? If not, problem is in the Telnet code. Otherwise, continue with Step 4.
4. Does VTAM buffer trace show data passed to host application? If not, problem is in VTAM code. If buffer trace shows correct data, continue with Step 5.
5. Does the buffer trace show data coming from the host application? If not, the problem is in the host application. Contact your host application support center for these products. Otherwise, continue with Step 6.
6. Does the buffer trace show data sent back to the Telnet LU? If not, the problem is in VTAM. Otherwise, continue with Step 7.
7. Is the last data from the application seen in the CTRACE with TELNET option output? If not, the problem is in Telnet. Otherwise, continue with Step 8.
8. Does the packet trace show the data sent to the client? If not, the error is in TCP/IP platform. Otherwise, continue with Step 9.
9. Check the data in the packet trace output to see whether unlock keyboard is set on in the data stream. If unlock is set in the output data, the problem is in the emulator or client code. Otherwise, continue with Step 10.
10. Check the last data received by the Telnet LU in the VTAM buffer trace. If unlock is set in that data stream, or end bracket or change direction is set in the RH, the problem is in the Telnet code. If none is set, the host application did not allow for unlocking of the keyboard. Contact your host application software support.

Results

If the preceding problem determination shows the error to be in the TCP/IP platform or Telnet code, a dump is needed to allow a more detailed investigation of the problem.

Incorrect output (server)

Problems with incorrect output are reported when the data sent to the client is not seen in its expected form. This could be garbled data that is unreadable on the screen, a blank screen when output is expected, or screen formatting problems. These problems are generally traced back to logmode issues. Ensure the primary and alternate screen sizes in the logmode used are correct for the TN3270 or TN3270E emulator that you are using. The logmode coded in the TCPIP profile is suggested to VTAM as the correct logmode for this device type. The VTAM PLU application determines the actual logmode that is used. Therefore this application must be configured correctly to use the appropriate logmode.

If a problem is recreatable, you can use the Telnet DEBUG features. See z/OS Communications Server: IP Configuration Guide for details.

Documentation

Documentation needed to find the source of the error in an incorrect output problem would be:

- CTRACE with TELNET option and the FULLDATATRACE parm active in the profile
- VTAM buffer trace of the Telnet LU, with AMOUNT=FULL specified
- Client screen output information

Steps for analyzing incorrect output (server)

The main goal of diagnosing this type of problem is to determine whether the data was sent incorrectly by the host application or corrupted by VTAM, TCP/IP, Telnet, or Telnet client code.

About this task

Table 33 lists the types of incorrect output that might be seen and the steps needed to identify the code in error.

Table 33. Incorrect output types for Telnet

Incorrect Output	Analysis Steps
Blank screen	1, 6, 7
Garbled or unreadable characters on the screen	2, 3, 4, 5, 6, 7
Incorrectly formatted screen	6, 7

Procedure

See Table 33 to identify which of the following steps to use in determining the cause of the error.

1. Was the last output data seen in a SEND DATA to CLIENT CTRACE entry displayed at the terminal emulator? If not, the problem is in the client or emulator. Contact your emulator provider for this product. If the last output

was seen at the terminal, go to step 9 on page 519 of the analysis procedure in "Session hangs (server)" on page 518, and continue your diagnosis.

2. Was the TELNET command entered with TRANSLATE specified? If so, make sure the translate table is compatible with the capabilities of the client device. If compatible or no TRANSLATE was used, continue with Step 4.
3. The CTRACE with TELNET option entries show the data as it arrived from VTAM and again as it goes to the stack. FULLDATATRACE parameter should be specified in the profile when looking for a problem in the data stream. Examine the CTRACE and compare the DATA from VTAM entries to the DATA to CLIENT entries. If they are different, then Telnet altered the data stream. If not, the problem is with the TCP/IP platform code. Otherwise, continue with Step 4 on page 519.
4. In the data trace output, is the data stream sent by the server the same as received from VTAM? If not, the problem is with the Telnet code. Otherwise, continue with Step 5 on page 519.

   Tip: If the client is an ASCII device, these might be different due to EBCDIC-to-ASCII translation. Check the appropriate translate table for compatibility with the client device.
5. In the VTAM Buffer trace with the FULL option specified, is the data in the VTAM USER entry (data received by VTAM) the same as the data in the VTAM BUFF entry (data sent by VTAM)? If not, VTAM has corrupted the data. Otherwise, incorrect data was sent by the application. Contact the IBM Software Support Center for the host application.
6. Is the LOGMODE specified for the negotiated terminal type valid for the actual client device?

   Tip: A VTAM session display specifying the SID for the session shows the actual logmode selected by the SNA application.
7. Does the device characteristics information in the BIND sent by the host application match the device characteristics information in the specified LOGMODE entry, and are these characteristics appropriate for the emulator in use? This can be checked by comparing the specified LOGMODE entry (see z/OS Communications Server: SNA Customization) with the BIND in the buffer trace at logon to the selected application. For information about the BIND RU, see z/OS Communications Server: SNA Programming and *SNA Formats*.

Results

If the problem is not found after using the analysis steps, contact your IBM Software Support Center for additional diagnostic suggestions.

Session outages (server)

Session outages are reported as an unexpected termination of the TCP/IP connection or the Telnet-to-host application session. A session that has been disconnected or terminated results in the client being returned to the panel where the initial TELNET command was entered and message EZZ6034I is issued. See z/OS Communications Server: IP Messages Volume 4 (EZZ, SNM).

Telnet sessions can be terminated due to TELNETPARMS specified in the PROFILE data set. Telnet ends a session if there is no activity on the SNA side of the connection for the amount of time specified in the INACTIVE parameter. Telnet checks for dormant sessions on the IP side of the connection using the

SCANINTERVAL/TIMEMARK parameters specified. When appropriate, the connection is terminated due to this processing. See the SCANINTERVAL/TIMEMARK parameters in z/OS Communications Server: IP Configuration Reference for additional information.

Documentation

The following documentation is needed for initial investigation of problems reported as session outages.

Abnormal connection terminations are reported by using the EZZ6034i message with the appropriate reason code (RCODE). If DEBUG SUMMARY is coded in the Telnet profile, then normal connection terminations are also reported. If DEBUG DETAIL is coded, then additional diagnostic information is reported by using EZZ6035I messages. These messages can be spooled to either the console or job log. Examination of the RCODE carried in these messages is the first step to diagnosing this type of problem.

Steps for analyzing session outages (server)

The preceding output is needed to begin diagnosis of a session outage reported against Telnet. It is also helpful to know what processing the Telnet user was doing at the time of the interrupted session.

Procedure

Perform the following steps:

1. If a timeout due to inactivity or termination due to TIMEMARK processing is suspected, check the values set in the PROFILE data set.
2. Additional messages are issued for session outages when the Telnet DEBUG features are active. See z/OS Communications Server: IP Configuration Guide for details of the Telnet DEBUG features.
3. Check the documentation listed in "Documentation" for indications of an error.
   - If the MVS system console indicates a VTAM error, continue diagnosis with your VTAM programmer.
   - If the console shows a Telnet or TCP/IP error, check z/OS Communications Server: IP Messages Volume 1 (EZA) or z/OS Communications Server: IP Messages Volume 4 (EZZ, SNM) and follow the directions for system programmer response for the message.

   If messages are found that do not lead to an accurate diagnosis and resolution of the error, search the APAR data base, available at for more information. If this does not provide a solution, contact the IBM Software Support Center.
4. If only one Telnet user session was affected, continue with step 5. Otherwise, go to step 7 on page 523.
5. If the problem can be re-created by performing the same operation or processing, run the following traces:
   - TCP/IP packet trace filtered using the IP address of the failing client
   - Component Trace output (CTRACE) specifying the Telnet option
   - VTAM Internal Trace (VIT)
   - VTAM buffer trace output with AMOUNT=FULL specified.

   Note: Contact your IBM Software Support Center for information about options needed before running these traces.
6. If all Telnet user sessions were interrupted, do one of the following:
   - Check the MVS system console and LOGREC for abends.

- Check for loss of network connectivity. Verify whether all the TELNET users come in through the same channel interface or through a common router.

7. If there are no messages or abends and all Telnet user sessions have been disconnected, the traces listed in Step 5 on page 522 are needed during a recurrence of the failure.

   A dump of the TCP/IP address space or a dump of the Telnet address space should be taken at this time. To capture the necessary areas of storage in the DUMP command, include:

   ```
   SDATA=(CSA,LSQA,PSA,RGN,SQA,SUM,SWA,TRT,LPA)
   ```

   Instructions on obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids for your release of MVS.

Special considerations when using SSL encryption support

Because data flowing across the connection between the client and the server is encrypted, the data field in the packet trace is also encrypted after SSL handshaking is completed. If problem determination requires seeing Telnet handshake or user data, you also need to run Component Trace to see the decrypted data field. When starting Component Trace, specify **options=**(*TELNET*) and use IPCS to format the Component Trace. For more information about Component Trace, see Chapter 5, "TCP/IP services traces and IPCS support," on page 45.

The Telnet Component Trace records contain the connection ID in the CID field. The connection ID in the trace corresponds to the connection ID output of the connection display command. Use this field to locate records related to the client in question. After an LUname has been assigned, the Component Trace User field shows the LUname, providing additional data for locating your client.

The following Component Trace records might be of interest:

SKSCINIT Succeeded
: SSL handshaking completed and subsequent data on this connection is encrypted.

Receive Data from Client
: The Data from Client field of this record contains the decrypted data coming from the client.

Send Data to Client
: The Data to Client field of this record contains the decrypted data going to the client.

Following is a sample Send Data to Client Component Trace record:

```
MVS181    TELNET    70010004  12:49:06.354966   Send Data to Client
HASID..002A        PASID...002A        SASID..002A       MODID..EZBTTSND
TCB....00000000    REG14...89D37F40    USER...TCPM1011   DUCB...0000000D
CID....092552C4    SEQ.....000024BE
  ...
  ...
  ADDR...00000000  08167AB0  LEN....00000004  Number of Bytes Sent
    +0000  0000002C                                 | ....            |
  ADDR...00000000  7F687950  LEN....0000002C  Data to Client
    +0000  F5C1115D  7F1D4011  40401DC8  C9D2D1F5  | 5A.)". .  .HIKJ5 |
    +0010  F6F7F0F0  C140C5D5  E3C5D940  E4E2C5D9  | 6700A ENTER USER |
    +0020  C9C44060  1D4011C1  5013FFEF            | ID -. .A&...     |
```

Telnet Component Trace data

To help associate a Component Trace entry with a particular client, the following two Component Trace fields contain data unique to Telnet:

CID The connection ID for the connection. This is equivalent to the connection ID output from the connection display command.

USER The LUname associated with the client, after it has been assigned. Prior to LUname assignment, this field might be null or contain the TCP procedure name. The LUname is not set until after the completion of the Telnet handshake.

Use these fields in Component Trace formatting to limit the records to be displayed. For example, if you want Telnet records for a client connection ID X'021F' with the LUName TCPM1011, code the following IPCS command:

```
CTRACE COMP(SYSTCPIP) SUB((proc_name)) FULL JOBLIST (TCPM1011)
OPTIONS((TELNET,CID(X'0000021F')))
```

Tip: Some of the records pertinent to the connection are not shown when the output is restricted by the CID and USER options. However, it is often helpful to use the output produced by these filters as a starting point.

General Telnet client information

The Telnet client code runs under TSO in the TSO user's address space. The Telnet client uses the VTAM interface, like other TSO applications, to send data out to the user's terminal.

The Telnet client can run in line mode, when accessing an ASCII host, or run in full-screen mode, if the remote host provides 3270 full-screen support.

Telnet client definitions

The Telnet command must be authorized to be issued by TSO users. See the z/OS MVS Initialization and Tuning Guide for information about making Telnet an authorized command. There are no other special definitions or setup requirements to run the Telnet client.

Diagnosing Telnet client problems

Problems that might involve the Telnet client are reported as one of the following types:

- Abends
- Session hangs
- Incorrect output

Use the information in the following topics for problem determination and diagnosis of errors reported in the Telnet client.

Abends (client)

An abend in the TELNET client should result in messages and error-related information being sent to the MVS system console. These abends should affect only the TSO user that was running Telnet. A dump of the error is needed unless the symptoms match a known problem.

Documentation

Code a SYSMDUMP DD or SYSABEND DD statement in the TSO PROC to ensure that a useful dump is obtained in the event of an abend. See Chapter 3, "Diagnosing abends, loops, and hangs," on page 21, for more information.

Analysis

See z/OS Problem Management or see Chapter 3, "Diagnosing abends, loops, and hangs," on page 21 for more information about debugging dumps produced during TCP/IP processing.

Session hangs (client)

This topic discusses diagnosis of a hang after a session has been successfully connected. A hang is indicated by the keyboard remaining locked after sending or receiving data from the remote host.

There are many components involved in the transfer of data from a locally attached device through a Telnet session. Any one of these might be the cause or a contributing factor to the hang. Each must be investigated to define the area responsible for the failure.

Documentation

To determine the cause of a Telnet client session hang, the following is needed:

- Information about what was seen at the client screen
- VTAM buffer trace of the local device LU
- VTAM internal trace (if the error appears to be in VTAM)
- VTAM TSO trace of the user ID issuing Telnet
- GTF trace of SVC93 and SVC94 (TGET/TPUT)
- Telnet client trace
- Dump of the TSO user's address space
- TCP/IP packet trace and CTRACE with TELNET option on remote host (if possible)

The preceding list of documentation is a complete list that includes documentation needed to resolve most types of hangs. All of the indicated data might not be needed for each occurrence of a hang. The following analysis section provides information about what types of data might be needed through each diagnostic step.

Steps for analyzing session hangs (client)

Diagnosing a Telnet client hang.

Before you begin

To assist with diagnosis of a Telnet client hang, it is helpful to be familiar with the components involved and understand which ones interface directly with each other. In the case of a Telnet from an MVS client to a remote host, the following occurs:

- Data is entered by the user and then passed by VTAM to TSO.
- Data is passed from TSO to Telnet client code.
- Data is transferred across the TCP/IP connection to the remote host.
- The remote server sends data to the target application.

Note: It is suggested that a VTAM buffer trace and a Telnet client trace be run while recreating the problem for initial debugging purposes. A sample of the client trace output can be found in Figure 61 on page 529. See z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures or *SNA Formats* for more information about VTAM buffer trace output.

Procedure

Perform the following steps:

1. Does the hang affect other Telnet clients? If so, go to "Diagnosing TN3270E Telnet server problems" on page 516. Otherwise, continue with Step 2.
2. Was the last activity at the terminal input or output? If input, go to step 5. If output, continue with Step 3.
3. Check the data in the VTAM buffer trace to see whether unlock keyboard is set on in the data stream. If unlock is set on in the data stream, the problem is in the emulator, control unit, or terminal device. If not, check the Telnet client trace to ensure the output data stream matches what is seen in the buffer trace. If the data streams match, the remote host application has not unlocked the keyboard. Contact your IBM Software Support Center for the host application for more help with the problem. If the data streams do not match, continue with Step 4.
4. The problem appears to be in the VTAM TSO area. Re-create the error while running the Telnet client trace, a GTF trace of SVC93 and SVC94, a VTAM TSO trace, and a VTAM buffer trace. Contact your IBM Software Support Center for assistance in interpreting the traces.
5. Check the VTAM buffer trace to ensure input data was received by VTAM and passed to TSO. If the last data entered at the terminal is not in the VTAM buffer trace, the problem is in the PC emulation code or in the control unit. If input data is correct, continue with Step 6.
6. Is the entered data seen in client trace output? If not, the problem is in VTAM TSO. Follow the instructions in Step 4. If data is in the client trace, the error needs to be diagnosed from the server host. See "Session hangs (server)" on page 518 and follow the path for "last activity at the terminal was input."

Results

Documentation listed earlier, but not referenced in the previous debugging steps, can be useful in the following situations:

- VTAM internal trace

  Note: Data is seen in "BUFF VTAM" VTAM buffer trace entry (entering VTAM from the terminal), but not in the "BUFF USER" VTAM buffer trace entry (passed from VTAM to TSO).
- Dump of TSO user's address space

  Note: Data is seen in the "BUFF USER" VTAM buffer trace entry, but not in the VTAM TSO trace or Telnet client trace.

Contact the IBM Software Support Center for assistance with further diagnosis when data is obtained in these situations.

Tip: Information about starting and examining traces is discussed in "Step for starting Telnet client traces" on page 528.

Incorrect output (client)

Problems with incorrect output are reported when the data seen at the terminal is not in its expected form. This might be garbled data that is unreadable, a blank screen when output is expected, or screen formatting problems.

Documentation

Documentation needed to find the source of the error in an incorrect output problem is:

- VTAM buffer trace of the local device LU
- VTAM TSO trace of the user ID issuing Telnet
- GTF trace of SVC93 and SVC94
- Telnet client trace
- Client screen output information

Steps for analyzing incorrect output (client)

The main goal of diagnosing this type of problem is to determine whether the data was sent incorrectly by the host application or corrupted by the Telnet server, Telnet client, TSO, or VTAM code. The following analysis steps hopefully allow quick determination of whether the problem is a Telnet client problem or must be addressed from the server host.

Procedure

1. If new data sent to the screen cannot be read (garbled or formatted incorrectly), go to step 4. Otherwise, continue with Step 2.
2. Was the last output data seen in the VTAM buffer trace displayed at the terminal? If not, the problem is in the emulator or device. Contact the appropriate IBM Software Support Center. Otherwise, continue with Step 3.
3. Does the last output data in the Telnet client trace match the data in the VTAM buffer trace? If not, contact your IBM Software Support Center with the client trace, a VTAM TSO trace, and a VTAM buffer trace of the error. Otherwise, this problem must be investigated from the Telnet server side. Continue with the investigation as a Telnet server session hang.
4. Was the TELNET command entered with TRANSLATE specified? If so, make sure the translate table is compatible with the capabilities of the output device. If the table is compatible or no TRANSLATE was used, continue with Step 5.
5. Check the Telnet client trace and VTAM buffer trace. If the data is different, contact your IBM Software Support Center with the client trace, a VTAM TSO trace, and a VTAM buffer trace. Otherwise, continue investigating as a Telnet server incorrect output problem.
6. If the data is formatted incorrectly for the screen size, check the defined session parameters for the negotiated device type for the Telnet server.

Results

If the problem is not found after using the analysis steps, contact your IBM Software Support Center for more diagnostic suggestions.

Telnet client traces

The Telnet client trace shows data received from the remote server to be sent to the local device, and data from the device to be forwarded to the remote host. This includes attention interrupts and some negotiation data seen at the beginning of the session. Data from the initial Telnet negotiation is not seen, only an indication that it is negotiation data and the number of bytes received.

Step for starting Telnet client traces

Before issuing the Telnet command, the following command should be issued from the TSO `ready` prompt or command line to allocate the trace data set:

```
ALLOC F(DEBUGFIL) DA(data.set.name) NEW
```

Trace data is written to the data set indicated in the command.

The trace is invoked by issuing the Telnet command with the DEBUG option:

```
TELNET hostname (DEBUG
```

Trace example (client)

Figure 61 on page 529 is sample output from a Telnet client trace showing part of a Telnet login to a remote host.

```
 1  EZA8310I DataDelivered; # bytes: 3
    EZA8338I ord: 255 asis:
    EZA8345I in TelnetRead
    EZA8305I in IacNoteArrives
 2  EZA8306I Option neg. stuff arrives
    EZA8310I DataDelivered; # bytes: 6
    EZA8338I ord: 255 asis:
    EZA8345I in TelnetRead
    EZA8305I in IacNoteArrives
    EZA8306I Option neg. stuff arrives
    EZA8310I DataDelivered; # bytes: 12
    EZA8338I ord: 255 asis:
    EZA8345I in TelnetRead
    EZA8305I in IacNoteArrives
    EZA8306I Option neg. stuff arrives
    EZA8338I ord: 255 asis:
    EZA8345I in TelnetRead
    EZA8305I in IacNoteArrives
    EZA8306I Option neg. stuff arrives
    EZA8338I ord: 255 asis:
    EZA8345I in TelnetRead
    EZA8305I in IacNoteArrives
    EZA8306I Option neg. stuff arrives
    EZA8338I ord: 255 asis:
    EZA8345I in TelnetRead
    EZA8305I in IacNoteArrives
    EZA8306I Option neg. stuff arrives
    EZA8310I DataDelivered; # bytes: 222
 3  EZA8359I Data received from TCP:
 4  EZA8361I FF FD 00 FF FB 00 05 C2 11 40 40
1D E4 C5 95 A3 85 99 40 E8
    EZA8361I 96 A4 99 40 E4 A2 85 99 89 84 7A 1D C4 00 00 00 00 00
00 00
    EZA8361I 00 1D E4 11 C1 50 1D E4 D7 81 A2 A2 A6 96 99 84 7A 1D
CC 00
    EZA8361I 00 00 00 00 00 00 00 1D E4 11 C1 F7 1D E4 D5 85 A6 40
97 81
    EZA8361I A2 A2 A6 96 99 84 7A 1D CC 00 00 00 00 00 00 00 00 1D
E4 11
    EZA8361I C2 60 1D E4 C1 97 97 93 89 83 81 A3 89 96 95 7A 1D C4
40 40
    EZA8361I 40 40 40 40 40 40 1D E4 11 C3 F0 1D E8 C1 97 97 93 89
83 81
    EZA8361I A3 89 96 95 40 99 85 98 A4 89 99 85 84 4B 40 D5 96 40
C9 95
    EZA8361I A2 A3 81 93 93 81 A3 89 96 95 40 C4 85 86 81 A4 93 A3
40 40
    EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
40 40
    EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 00 00 00 11 40 40
05 13
    EZA8361I FF EF
 5  EZA8364I z" w" B"  "UEnter Your Userid:D
" "" """U"A&"UPassword:">";
    EZA8364I "U"A7"UNew password:">"U"B-"UApplication:"D
    EZA8364I       "U"C0"YApplication required. No Installation Default
    EZA8364I                                  "   "┬Q
```

Figure 61. Telnet client trace

```
    EZA8339I In Transparent mode, found IAC at IacOffset 0, CurrentChar
is 0
    EZA8345I in TelnetRead
    EZA8305I in IacNoteArrives
    EZA8306I Option neg. stuff arrives
    EZA8339I In Transparent mode, found IAC at IacOffset 0, CurrentChar
```

```
is 3
     EZA8345I in TelnetRead
     EZA8305I in IacNoteArrives
     EZA8306I Option neg. stuff arrives
     EZA8339I In Transparent mode, found IAC at IacOffset 214, CurrentChar
is 6
     EZA8345I in TelnetRead
 6   EZA8313I got USERdeliversLINE
     EZA8371I in SendData
 7   EZA8380I User data is...
     EZA8381I 7D '
     EZA8381I C2 B
     EZA8381I F1 1
     EZA8381I 11 "
     EZA8381I 40
     EZA8381I D4 M
     EZA8381I E4 U
     EZA8381I E2 S
     EZA8381I C5 E
     EZA8381I D9 R
     EZA8381I F2 2
     EZA8381I 11 "
     EZA8381I C2 B
     EZA8381I 6E >
     EZA8381I E3 T
     EZA8381I E2 S
     EZA8381I D6 O
     EZA8381I 40
     EZA8381I 40
     EZA8381I 40
     EZA8381I 40
     EZA8381I 40
 8   EZA8382I ; Len is 22
     EZA8310I DataDelivered; # bytes: 48
     EZA8359I Data received from TCP:
     EZA8361I 05 C1 11 5D 7F 1D 40 11 40 40 1D C8 C9 D2 D1 F5 F6 F7
F0 F0
     EZA8361I C1 40 C5 D5 E3 C5 D9 40 E4 E2 C5 D9 C9 C4 40 60 1D 40
11 C1
     EZA8361I 50 13 FF EF 01 C2 FF EF
     EZA8364I  A")"" "  "HIKJ56700A ENTER USERID -" "A&"┬Q"B"Q;
     EZA8339I In Transparent mode, found IAC at IacOffset 42, CurrentChar
is 0
     EZA8345I in TelnetRead
     EZA8339I In Transparent mode, found IAC at IacOffset 2, CurrentChar
is 44
     EZA8345I in TelnetRead
     EZA8313I got USERdeliversLINE
 9   EZA8371I in SendData
     EZA8380I User data is...
     EZA8381I 7D '
     EZA8381I C1 A
     EZA8381I D5 N
     EZA8381I 11 "
     EZA8381I 40
     EZA8381I 5A !
     EZA8381I A4 u
     EZA8381I A2 s
     EZA8381I 85 e
     EZA8381I 99 r
     EZA8381I F3 3
     EZA8382I ; Len is 11
     EZA8310I DataDelivered; # bytes: 1106
     EZA8359I Data received from TCP:
     EZA8361I 05 C3 11 40 40 3C 40 40 40 11 40 40 1D E8 60 60 60 60 60 60
     EZA8361I 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60
     EZA8361I 60 60 60 60 60 40 E3 E2 D6 61 C5 40 D3 D6 C7 D6 D5 40 60 60
```

```
EZA8361I 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60
EZA8361I 60 60 60 60 60 60 60 60 60 60 60 60 60 60 11 C1 50 1D E8 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 11
EZA8361I C2 60 1D E8 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 11 5B 60 1D E8 D7 C6 F1 61 D7 C6 F1 F3 40 7E 7E
EZA8361I 6E 40 C8 85 93 97 40 40 40 40 D7 C6 F3 61 D7 C6 F1 F5 40 7E
EZA8361I 7E 6E 40 D3 96 87 96 86 86 40 40 40 40 D7 C1 F1 40 7E 7E 6E
EZA8361I 40 C1 A3 A3 85 95 A3 89 96 95 40 40 40 40 D7 C1 F2 40 7E 7E
EZA8361I 6E 40 D9 85 A2 88 96 A6 11 5C F0 1D E8 E8 96 A4 40 94 81 A8
EZA8361I 40 99 85 98 A4 85 A2 A3 40 A2 97 85 83 89 86 89 83 40 88 85
EZA8361I 93 97 40 89 95 86 96 99 94 81 A3 89 96 95 40 82 A8 40 85 95
EZA8361I A3 85 99 89 95 87 40 81 40 7D 6F 7D 40 89 95 40 81 95 A8 40
EZA8361I 85 95 A3 99 A8 40 86 89 85 93 84 11 C3 F3 1D E8 C5 95 A3 85
EZA8361I 99 40 D3 D6 C7 D6 D5 40 97 81 99 81 94 85 A3 85 99 A2 40 82
EZA8361I 85 93 96 A6 7A 11 C4 E3 1D E8 D9 C1 C3 C6 40 D3 D6 C7 D6 D5
EZA8361I 40 97 81 99 81 94 85 A3 85 99 A2 7A 11 C6 D2 1D 60 40 E4 A2
EZA8361I 85 99 89 84 40 40 40 40 7E 7E 7E 6E 11 C6 E2 1D E8 E4 E2 C5
EZA8361I D9 F3 40 40 1D F0 11 C8 F2 1D 60 40 D7 81 A2 A2 A6 96 99 84
EZA8361I 40 40 7E 7E 7E 6E 11 C9 C2 1D 4C 00 00 00 00 00 00 00 00 1D
EZA8361I F0 11 4D F2 1D 60 40 C1 83 83 A3 40 D5 94 82 99 40 7E 7E 7E
EZA8361I 6E 11 4E C2 1D C8 00 00 00 00 00 00 00 00 00 00 00 00 00 00
EZA8361I 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
EZA8361I 00 00 00 00 00 00 1D F0 11 4B D2 1D 60 40 D7 99 96 83 85 84
EZA8361I A4 99 85 40 7E 7E 7E 6E 11 4B E2 1D C8 D4 E5 E2 F4 F2 F2 40
EZA8361I 40 1D F0 11 50 D2 1D 60 40 E2 89 A9 85 40 40 40 40 40 40 7E
EZA8361I 7E 7E 6E 11 50 E2 1D C8 F4 F0 F9 F6 00 00 00 1D F0 11 D2 F2
EZA8361I 1D 60 40 D7 85 99 86 96 99 94 40 40 40 7E 7E 7E 6E 11 D3 C2
EZA8361I 1D C8 00 00 00 1D F0 11 4C C2 1D 60 40 C7 99 96 A4 97 40 C9
EZA8361I 84 85 95 A3 40 40 7E 7E 7E 6E 11 4C D5 1D C8 00 00 00 00 00
EZA8361I 00 00 00 1D F0 11 C9 E2 1D 60 40 D5 85 A6 40 D7 81 A2 A2 A6
EZA8361I 96 99 84 40 7E 7E 7E 6E 11 C9 F5 1D 4C 00 00 00 00 00 00 00
EZA8361I 00 1D F0 11 D7 F3 1D E8 C5 95 A3 85 99 40 81 95 40 7D E2 7D
EZA8361I 40 82 85 86 96 99 85 40 85 81 83 88 40 96 97 A3 89 96 95 40
EZA8361I 84 85 A2 89 99 85 84 40 82 85 93 96 A6 7A 1D 60 11 D9 C7 1D
EZA8361I E8 00 11 D9 C9 1D C8 40 1D F0 60 D5 96 94 81 89 93 1D 60 11
EZA8361I D9 D7 1D E8 00 11 D9 D9 1D C8 40 1D F0 60 D5 96 95 96 A3 89
EZA8361I 83 85 1D 60 11 D9 E8 1D E8 00 11 D9 6A 1D C8 00 1D F0 60 D9
EZA8361I 85 83 96 95 95 85 83 A3 1D 60 11 D9 7A 1D E8 00 11 D9 7C 1D
EZA8361I C8 40 1D F0 60 D6 C9 C4 83 81 99 84 40 1D 60 11 D5 D2 1D 60
EZA8361I 40 C3 96 94 94 81 95 84 40 40 40 7E 7E 7E 6E 11 D5 E2 1D C8
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40 40
EZA8361I 1D F0 11 C7 C2 1D 7C 40 E2 85 83 93 81 82 85 93 40 40 40 40
EZA8361I 40 7E 7E 7E 6E 11 C7 D5 1D 7C 40 40 40 40 40 40 40 40 1D F0
EZA8361I 11 C9 C3 13 FF EF
EZA8364I C -    Y -------------------------------- TSO/E LOGON --
EZA8364I ------------------ A&  Y
EZA8364I
EZA8364I B- Y
EZA8364I                          |$- YPF1/PF13==> Help PF3/PF15=
EZA8364I Logoff    PA1==> Attention PA2==> Reshow *0 YYou may
EZA8364I request specific help information by entering a '?' in any
EZA8364I entry field C3 YEnter LOGON parameters below: DT YRACF LOGON
EZA8364I parameters: FK - Userid     ===> FS YUSER3  0 H2 - Password
EZA8364I ===> IB <           0 (2 - Acct Nmbr ===> +B H
EZA8364I 0".K - Procedure===>  .S  HMVS422   EZA8364I 0
      &K - Size      ===>  &S  H4096   0 K2 - Perform   ===> LB
EZA8364I H    0 <B - Group Ident  ===>  <N H 0 IS- New Passw
EZA8364I ord ===> I5 < 0 P3 YEnter an 'S' before each option
```

```
EZA8364I desired below: - RG Y  RI H 0-Nomail RP Y  RR H  0-Nonoti
EZA8364I ce - RY Y R: H  0-Reconnect - R: Y  R@ H  0-OIDcard - NK -
EZA8364I  Command   ===> NS H
EZA8364I                                          0 GB @ Seclabel
EZA8364I  ===>GN @        0 IC 'Q
EZA8339I In Transparent mode, found IAC at IacOffset 1104, CurrentChar is 0
EZA8345I in TelnetRead
EZA8313I got USERdeliversLINE
EZA8371I in SendData
```

Following are short descriptions of the numbered items in the trace:

1 This entry shows the data received from the Telnet server and indicates the number of bytes. The example here is during initial negotiation and does not include the actual data received.

2 This indicates the type of data received.

3 This entry indicates the data received from TCP (from the Telnet server).

4 The actual hexadecimal data received. This trace example is of a transparent mode session, so the data is in EBCDIC. In a line mode session, the data would be in ASCII, and there would be one character per line (like the input data later in the trace).

5 This is the translation of the previous hexadecimal data. All hexadecimal characters that translate into readable data are displayed.

6 This entry indicates data received from the terminal or PC.

7 Following this line is the actual input data. There is a single hexadecimal byte per line that is translated into its readable form.

8 This entry follows the input data and indicates the number of bytes received from the terminal.

9 This entry indicates the data from the host application (using the Telnet server) that is being sent to the terminal.

Telnet commands and options

For information about Telnet connection negotiations, see RFC 2355. Table 34 describes the Telnet commands from RFC 854, when the codes and code sequences are preceded by an IAC. For more information about Telnet commands, see RFC 854.

Table 34. Telnet commands from RFC 854

Command	Code	Description
SE	X'F0'	End of subnegotiation parameters.
NOP	X'F1'	No operation.
Data Mark	X'F2'	The data stream portion of a Synch. This should always be accompanied by a TCP Urgent notification.
Break	X'F3'	NVT character BRK.
Interrupt Process	X'F4'	The function IP.
Abort output	X'F5'	The function AO.
Are You There	X'F6'	The function AYT.
Erase character	X'F7'	The function EC.
Erase Line	X'F8'	The function EL.

Table 34. Telnet commands from RFC 854 (continued)

Command	Code	Description
Go ahead	X'F9'	The GA signal.
SB	X'FA'	Indicates that what follows is subnegotiation of the indicated option.
WILL (option code)	X'FB'	Indicates the want to begin performing, or confirmation that you are now performing, the indicated option.
WON'T (option code)	X'FC'	Indicates the refusal to perform, or continue performing, the indicated option.
DO (option code)	X'FD'	Indicates the request that the other party perform, or confirmation that you are expecting the other party to perform, the indicated option.
DON'T (option code)	X'FE'	Indicates the demand that the other party stop performing, or confirmation that you are no longer expecting the other party to perform, the indicated option.
IAC	X'FF'	Data byte 255.

Table 35 lists the options available for Telnet commands from RFC 1060. For more information about Telnet protocols, see RFC 1060 and RFC 1011.

Table 35. Telnet command options from RFC 1060

Option	Option (Hex)	Name
0	0	Binary Transmission
1	1	Echo
2	2	Reconnection
3	3	Suppress Go Ahead
4	4	Approx Message Size Negotiation
5	5	Status
6	6	Timing Mark
7	7	Remote Controlled Trans and Echo
8	8	Output Line Width
9	9	Output Page Size
10	A	Output Carriage-Return Disposition
11	B	Output Horizontal Tab Stops
12	C	Output Horizontal Tab Disposition
13	D	Output Formfeed Disposition
14	E	Output Vertical Tabstops
15	F	Output Vertical Tab Disposition
16	10	Output Linefeed Disposition
17	11	Extended ASCII
18	12	Logout
19	13	Byte Macro
20	14	Data Entry Terminal

Table 35. Telnet command options from RFC 1060 (continued)

Option	Option (Hex)	Name
21	15	SUPDUP
22	16	SUPDUP Output
23	17	Send Location
24	18	Terminal Type
25	19	End of Record
26	1A	TACACS User Identification
27	1B	Output Marking
28	1C	Terminal Location Number
29	1D	Telnet 3270 Regime
30	1E	X.3 PAD
31	1F	Negotiate About Window Size
32	20	Terminal Speed
33	21	Remote Flow Control
34	22	Linemode
35	23	X Display Location
255	FF	Extended-Options-List

Chapter 17. Diagnosing Simple Mail Transfer Protocol (SMTP) problems

The Simple Mail Transfer Protocol (SMTP) is used to transfer electronic mail reliably and efficiently. Recipients of the mail can be users on a local host, users on Network Job Entry (NJE), or users on remote TCP/IP hosts. The SMTPNOTE command is used to the send mail to a local or remote host.

This topic describes how to diagnose problems with SMTP and contains the following subtopics:

- "Sender SMTP"
- "Receiver SMTP"
- "SMTP environment"
- "SMTP definitions" on page 536
- "Diagnosing SMTP problems" on page 536
- "ADDRBLOK data set" on page 542
- "SMTP RESOLVER trace" on page 544

For information about diagnosing problems with the other z/OS Communications Server mail application, z/OS UNIX sendmail, see Chapter 18, "Diagnosing z/OS UNIX sendmail and popper problems," on page 547.

Sender SMTP

The sender SMTP performs the following functions:

- Receives notes from the SMTPNOTE CLIST by way of a TSO TRANSMIT command or through a batch job using IEBGENER
- Resolves the host name of recipients by way of the RESOLVER module
- Opens a TCP/IP connection with the SMTP server
- Returns mail to the sender, if mail is undeliverable

Receiver SMTP

The receiver SMTP provides the following functions:

- Accepts mail from remote TCP/IP hosts
- Delivers mail to the local user using TSO TRANSMIT to the spool for the local user
- Forwards mail to the next "hop", if this is not the final destination
- Rejects mail for recipients who are not valid

SMTP environment

Figure 62 on page 536 shows the SMTP environment.

© Copyright IBM Corp. 2000, 2015

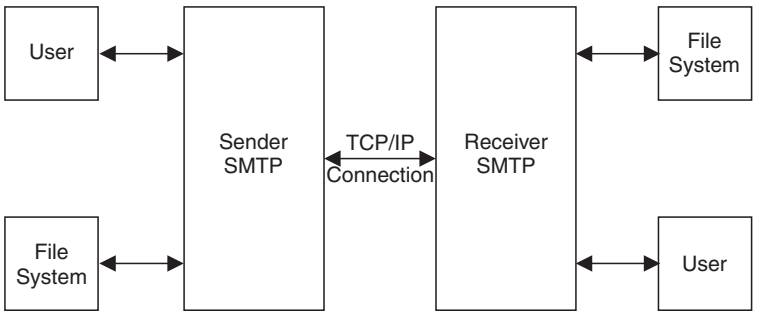

Figure 62. SMTP environment

SMTP definitions

In order to run correctly, SMTP must be defined correctly for both TCP/IP and SMTP. The SMTP.CONFIG and TCPIP.DATA data sets contain the main sender and receiver parameters. The SMTPNOTE CLIST must be customized for your particular installation. The *nodename* must have been specified during VMCF initialization. If you have configured VMCF and TNF as non-restartable subsystems, the *nodename* is specified in the IEFSSNxx member of PARMLIB. If you have configured VMCF and TNF as restartable subsystems, the *nodename* is specified as the value of the P= parameter of the EZAZSSI started procedure.

For more information about starting VMCF and TNF, see z/OS Communications Server: IP Configuration Guide.

Restrictions:

- The NJE node name, *nodename*, must be the same as the *hostname* and the *smtpnode* in the SMTPNOTE CLIST.
- SMTP can handle only one NJE node name.

See z/OS Communications Server: IP Configuration Reference for more information about configuring SMTP.

Diagnosing SMTP problems

SMTP problems are generally reported under one of the following categories:

- Abend
- Spooling
- SMTP does not deliver mail
- SMTP loop
- Mail item has incorrect output

Abends

An abend during SMTP processing should result in messages and error-related information being sent to the system console. A dump of the error is needed unless the symptoms already match a known problem.

Documentation

The following documentation is needed for abends:

- Dump

Guideline: Code a SYSMDUMP DD or SYSABEND DD statement in the SMTP cataloged procedure to ensure that a useful dump is obtained in the event of an abend.

- Output from the started SMTP procedure
- SYSLOG and LOGREC output for the time of the error

Analysis

See z/OS Problem Management or see Chapter 3, "Diagnosing abends, loops, and hangs," on page 21, for information about debugging dumps produced during SMTP processing.

Spooling problems

Spooling problems can occur when the VERB command is being used and the origination information is either missing or not valid. The VERB command requires the originator to have a valid JES user ID and node ID on the SMTP sending system. The originator information is taken from the TSO XMIT (Transmit) command headers.

For more information about the VERB command, see z/OS Communications Server: IP User's Guide and Commands.

SMTP does not deliver mail

This topic discusses diagnosis of mail items that are not delivered to the recipient. Problems with mail not being forwarded can be divided into the following categories:

- Mail not forwarded to a local user
- Mail not forwarded to a user on another NJE host
- Mail not forwarded to remote TCP/IP host

Steps for undeliverable mail items

These steps apply to all categories of this problem.

Procedure

Perform the following steps:

1. If an undeliverable mail message is returned to the originator of the mail, obtain the undeliverable mail message from the originator. It contains text indicating why the mail could not be sent.
2. If an undeliverable mail message is not returned to the originator of the mail.
   a. Check the SMTP log to see whether the undeliverable mail message is sent to the postmaster.
      - Check whether an SMTP EXIT program is installed and activated for outbound mail.
      - Check the SMTP.CONFIG data set for the EXITDIRECTION BOTH statement.
      - If EXITDIRECTION BOTH is coded, activate DEBUG in SMTP.CONFIG data set.
      - Check SYSDEBUG log to see whether SMTP exit program is rejecting the mail.
      - If yes, check the SMTP exit program.
   b. Check the SMTP log to see whether the undeliverable mail message is deleted because the mailbox of the originator does not exist.

Documentation

The following documentation should be available for initial problem diagnosis:

- TSO console log with the SMTPNOTE messages
- Job log output from the started SMTP procedure
- SMTP.CONFIG data set
- TCPIP.DATA data set

Other documentation that might be needed is discussed in the following topic.

Steps for analyzing mail delivery problems

Analyze why mail is not delivered properly.

Procedure

Perform the following steps:

- If the problem is that mail was not forwarded to a local user:
  1. Was SMTPNOTE customized for your installation?
  2. Is the local user one that is coded as a restricted user in the SMTP.CONFIG data set?
  3. Are the JES node parameters coded correctly? This can be determined by issuing a TSO TRANSMIT of a data set to the user and node. If the transmission works, the JES node parameters are coded correctly.
  4. Activate DEBUG in SMTP.CONFIG data set. Check SYSDEBUG log to see whether SMTP exit program is rejecting the mail. If yes, customer needs to check SMTP exit program.
  5. If TSO TRANSMIT fails with message `INMX202I Node name SMTPNODE not defined to JES` when testing customization of SMTPNOTE variables, check that the SMTPNODE variable used by SMTPNOTE is defined correctly in the JES2PARM data set as a node name. Also check that the SMTPJOB name used by SMTPNOTE is not defined as a node name to JES.
- If the problem is that a mail note was not forwarded to an NJE host:
  1. Follow the preceding steps for mail that was not forwarded to a local user.
  2. Is SMTP configured as an NJE gateway?
  3. Was SMTPNJE successfully run to create the NJE host table data set?
  4. Check whether the NJE node is in the NJE host table data set.

  See the z/OS Communications Server: IP Configuration Reference for information about SMTP configuration.
- If the problem is that mail was not forwarded to a remote TCP/IP host:
  1. Determine whether the mail is queued on the resolver queues or on a specific IP address queue.
     a. Issue the **queue** command by doing one of the following tasks:
        1) Issue the **modify** command. See z/OS Communications Server: IP System Administrator's Commands for more information about the MODIFY command for SMTP.
        2) Issue the SMSG SMTP QUEUE command on the TSO command line.
     b. Browse the ADDRBLOK data set for obvious errors. The ADDRBLOK data set is described in "ADDRBLOK data set" on page 542.

Restriction: Stop SMTP to obtain the ADDRBLOK data set as it was sent, because the data set is updated during processing and deleted when the number of recipients equals 0.

2. Mail is on the resolver queues if the host name cannot be resolved to an IP address.
   a. Check the system console or the SMTP output log for an EZA5208E message. If this message is present, then SMTP has received a domain name server (DNS) response that has the truncation bit set to on in the packet. SMTP discards these packets because the DNS response is not complete. SMTP continues to retry these DNS queries for the amount of time that is configured. If the same DNS response continues, the mail message eventually becomes undeliverable. Use the RESOLVER traces to determine which DNS is sending the response and the unresolved name. Look for message EZA5526I in the resolver traces. Update the BIND server data base files or the TCPIP.DATA file to specify a correct set of name servers to resolve this problem.

      Tip: You can avoid this problem by migrating to the new CSSMTP application. The CSSMTP application handles the DNS truncation differently because it uses the system resolver. SMTP has its own resolution code.
   b. If it has not, run RESOLVER trace to see whether the host name is resolved correctly. The RESOLVER trace is explained in "SMTP RESOLVER trace" on page 544.
3. Mail is on a specific IP address queue. Is this the expected IP address of the target mail server? If your answer is no, answer the following questions:
   – Are you using IPMAILERADDRESS? If you are, understand the implications of the following settings.

     RESOLVERUSAGE NO
     This setting causes all mail destination for IP networks to be sent to this IP address. The IP mailer address is wrong. If you recently changed the IP address, for this change to affect existing mail, issue the SMSG SMTP EXPIRE command on the TSO command line or see z/OS Communications Server: IP System Administrator's Commands for more information about the MODIFY command for SMTP.

     RESOLVERUSAGE YES
     This setting causes SMTP to send mail to the IPMAILERADDRESS only if it is unable to resolve the host name of the recipient of the mail. If SMTP has access to a domain name server, it will likely be able to resolve the host name and will not send mail to the IPMAILERADDRESS.

     If you want all mail to be sent to the IPMAILERADDRESS change this setting to NO. For this change to affect existing mail, after you set it to NO, issue the SMSG SMTP EXPIRE command on the TSO command line or see z/OS Communications Server: IP System Administrator's Commands for more information about the MODIFY command for SMTP.
   – Are you using IPMAILERNAME? If the IP address is not correct, run RESOLVER trace to see whether the host name is resolved correctly. The RESOLVER trace is explained in "SMTP RESOLVER trace" on page 544.

     Is message EZA5647E generated in the SMTP logs? Or when resolver traces are active, is the `Potential loop IP mailer =` trace message generated? If yes, HOME IP addresses in the IP list are associated with the

IPMAILERNAME. Activate resolver traces in SMTP to understand how SMTP resolved the IPMAILERNAME. Either correct the IPMAILERNAME, or remove the HOME IP address from the list of addresses that are associated with the IPMAILERNAME in the DNS database or local host tables.

4. Mail is on a specific IP address queue with the expected IP address of the target mail server. Is the remote TCP/IP/SMTP server running?

   Use the PING command to see whether the remote TCP/IP is running, or try using Telnet to access the IP address of the remote mail server using port 25. You can also use packet trace to ensure that SMTP can connect to this IP address and that another remote SMTP mailer is at this address.

   Guideline: Options that are coded in the SMTP.CONFIG data set directly affect how and when names are resolved by name servers, how often mail delivery is attempted, if there is a problem in the network or the remote NAME server, or if the SMTP server is not running.

What to do next

If the problem still occurs after following this procedure and making any needed changes and corrections, obtain the following documentation and contact the IBM Software Support Center:

- SMTP.CONFIG data set
- TCPIP.DATA data set
- Output from SYSERR and SYSDEBUG of the started SMTP procedure with DEBUG turned on
- ADDRBLOK data set

SMTP loop

This topic discusses diagnosis of the SMTP address space looping during processing.

Documentation

If SMTP is looping and printing out AMPX... messages to SYSERR, do the following:

- Examine the SYSERR output for AMPX... error messages and traceback information of called routines.
- Call the IBM Software Support Center with this information.

Tip: Coding the NOSPIE runtime parameter in the SMTP cataloged procedure might help alleviate a Pascal error recovery loop. For example, code:

```
//SMTP PROC MODULE=SMTP,DEBUG=,PARMS='NOSPIE',SYSERR=SYSERR
```

See Chapter 3, “Diagnosing abends, loops, and hangs,” on page 21 for more diagnostic information about diagnosing loops.

Mail item has incorrect output

Problems with incorrect output are reported when the recipient does not see the mail item in its expected form.

Documentation

Use the following documentation to confirm the source of the error:

- SMTP.CONFIG data set

- TCPIP.DATA data set
- Output from SYSERR and SYSDEBUG from the started SMTP procedure with DEBUG turned on
- A packet trace from TCP/IP and network trace facility output

  This documentation might be needed in cases where the actual data in the TCP/IP packets needs to be examined
- SMTPhlq.*.ADDRBLOK data set is a control file for SMTP processing.

  Note: You should stop SMTP to obtain the ADDRBLOK data set as it was sent, because the data set is updated during processing and deleted when the number of recipients equals zero.
- SMTPhlq.*.NOTE data set is the contents of the note being sent across the TCP/IP connection containing both headers and mail body.

Steps for analyzing incorrect mail output

The main goal in diagnosing an incorrect output problem is to determine where the corruption occurs. Is the data corrupted in SMTP, TCP/IP, or by something or someone on the network?

Procedure

Perform the following steps to analyze the problem:

- If the problem is that the received mail item has incorrect output:
  1. Is the correct translation table being used or could it have been customized to cause the error?

     Correct the translation error.
  2. Do TCP/IP and SMTP receive the correct output from the remote host?

     Obtain TCP/IP packet trace output or network trace facility output or both to see the actual data in the packets from the remote host.
  3. Analyze the output from SMTP DEBUG for obvious errors. The body of the note (mail item) is not shown in this output.
- If the problem is that the sent mail item has incorrect output:
  1. Is the correct translation table being used, or could it have been customized to cause the error?

     Correct the translation error.
  2. Was the correct data sent from SMTP or TCP/IP?

     Obtain a TCP/IP packet trace to see the actual data in the packets as they leave TCP/IP.
  3. Analyze the output from SMTP DEBUG for obvious errors. The body of the note (mail item) is not shown in this output.

What to do next

If the problem cannot be corrected by this procedure, and you believe that the problem is caused by either SMTP or TCP/IP, call the IBM Software Support Center for further diagnosis.

Forcing resolution of queued mail

Normally, the SMTP server resolves the MX or A records of a piece of mail and stores the mail in the data sets pointed to by the MAILFILEDSPREFIX keyword in

the SMTP configuration data set. If the mail cannot be delivered for some period of time, the IP addresses in the mail can become old or obsolete. The data set names for each piece of mail are:

mailfiledsprefix.number.ADDRBLOK
mailfiledsprefix.number.NOTE

There are two ways to force the SMTP server to resolve the addresses:

- The preferred method is to issue the SMSG SMTP EXPIRE command. See z/OS Communications Server: IP User's Guide and Commands and z/OS Communications Server: IP System Administrator's Commands for more information about this command.
- An alternate method is to modify the ADDRBLOK data set for the piece of mail. For each recipient record (records three through the end of the data set), if the first character of the record is an S, then change the S to an E, for expired. This causes SMTP to resolve that record in the ADDRBLOK data set the next time the SMTP server is started. To modify the ADDRBLOK data set, the data set must be zapped, or a local utility program must be used. The data set cannot be modified using the ISPF editor or IEBUPDATE.

ADDRBLOK data set

An ADDRBLOK data set is the master control file for SMTP and is used for tracking the status of a mail item during mail delivery. One ADDRBLOK data set is allocated for each piece of mail and is built when the mail is received. The data set is allocated with a high-level qualifier of MAILFILEDSPREFIX from the SMTP.CONFIG data set. The data set is updated during mail processing and is deleted when the number of recipients equals zero.

Guideline: You might need to stop SMTP to obtain the ADDRBLOK data set as it was sent, because the data set is updated during processing and deleted when the number of recipients equals zero.

Table 36 shows the format of Record 1 (the master control record) of an SMTP ADDRBLOK data set.

Table 36. Format of Record 1 of an SMTP ADDRBLOK data set

Characters	Description	Length (in characters)
1–7	Total number of recipients	7
8–14	Number of unresolved recipients	7
15–21	Number of recipients left to send this mail item to	7
22	Unused	1
23–30	File name of note file	8
31	Unused	1
32–39	Date	8
40	Unused	1
41–48	Time	8
49	Unused	1
50–53	Unused	4
54–55	Unused	2

Table 36. Format of Record 1 of an SMTP ADDRBLOK data set (continued)

Characters	Description	Length (in characters)
56	Key **Value** **Meaning** **B** BSMTP RPLY file **S** Spool file **M** Spool file from Mailer **T** File from TCP **E** Error file	1
Note: Characters 57–80 are optional data used only when the key (Character 56) is "S" or "M."		
57–64	Tag user ID	8
65–72	Tag node ID	8
73–80	Spool ID on the current system	8
77–80	Spool ID of the file source	4

Table 37 shows the format of Record 2 (for an unresolved From record) of an SMTP ADDRBLOK data set.

Table 37. Format of Record 2 (for an unresolved from record) of an SMTP ADDRBLOK data set

Characters	Description	Length (in characters)
1	Key **Value** **Meaning** **U** Unresolved	1
2	Sender path length (user host.domain)	1
3–4	Length of sender ID	2
5–(L1+4)	Sender ID (who sent the mail)	L1
(L1+5) –(L1+6)	Length of sender host.domain	2
(L1+7) –(L1+L2+6)	Sending host.domain	L2
(L1+L2+7)	Length of sender ID	1
(L1+L2+8) –(L1+L2+L3+7)	Sender ID (who sent the mail)	L3

Table 38 shows the format of Record 2 (for a resolved From record) of an SMTP ADDRBLOK data set.

Table 38. Format of Record 2 (for a resolved from record) of an SMTP ADDRBLOK data set

Characters	Description	Length (in characters)
1	Key **Value** **Meaning** **M** Resolved	1
2	Sender path length (user host.domain)	1
3–4	Length of sender ID	2

Table 38. Format of Record 2 (for a resolved from record) of an SMTP ADDRBLOK data set (continued)

Characters	Description	Length (in characters)
5–(L1+4)	Sender ID (who sent the mail)	L1
(L1+5) –(L1+6)	Length of sender host.domain	2
(L1+7) –(L1+L2+6)	Sending host.domain	(L1+L2+7)
(L1+L2+8)	Length of sender ID	1
(L1+L2+9) –(L1+L2+L3+8)	Sender ID (who sent the mail)	L3
(L1+L2+L3+9)	Length of encoded return path	1
(L1+L2+L3+10) –(L1+L2+L3+L4+9)	Encoded return path	L4

Table 39 shows the format of Records 3–*n* of an SMTP ADDRBLOK data set.

Table 39. Format of Record 3 (for an unresolved from record) of an SMTP ADDRBLOK data set

Characters	Description	Length (in characters)
1	Key **Value** **Meaning** **U** Unresolved **M** Resolved	1
2–5	Time-to-Live (TTL)	4
6	Length of return path	1
7–8	Length of recipient user ID	2
9–(L1+8)	Recipient user ID	L1
(L1+9) –(L1+11)	Length of recipient host.domain	2
(L1+12) –(L1+L2+11)	Recipient's host.domain	L2
(L1+L2+12)	Length of recipient path	1
(L1+L2+13) –(L1+L2+L3+12)	Recipient path	L3
(L1+L2+L3+13)	Number of IP addresses	1
(L1+L2+L3+14) –(L1+L2+L3+17)	IP address 1	4
Note: There can be up to 16 IP addresses listed.		

SMTP RESOLVER trace

The RESOLVER trace shows requests and responses sent to and received from name servers. It also shows if local hosts tables are used for name resolution. This trace helps you diagnose problems with host name resolution.

RESOLVER trace output from SMTP is included in the job log output from the started SMTP procedure.

Figure 63 on page 545 shows an example of RESOLVER trace output. Short descriptions of the numbered items in the trace follow the figure.

```
   Userid of Caller:         ARMSTRNG
   TCP Host Name:            RALVMFE1
   Domain Origin:            RALEIGH.IBM.COM
   Jobname of TCP/IP:        TCPCS
   Communicate Via:          UDP
   OpenTimeOut:              30
   MaxRetrys:                1
   NSPort:                   53
   NameServer Userid:        NAMESRV
1  NSInternetAddress(.1.) := 9.67.1.5
   NSInternetAddress(.2.) := 9.67.5.44
   Data set prefix used:     TCPCS.BTA1

   Resolving Name:           RICKA
   Result from InitResolver: OK
   Building Name Server Query:
   * * * * * Beginning of Message * * * * *
2  Query Id:                 1
3  Flags:                    0000 0001 0000 0000
   Number of Question RRs:   1
4  Question   1: RICKA.RALEIGH.IBM.COM A IN
   Number of Answer RRs:     0
   Number of Authority RRs:  0
   Number of Additional RRs: 0
   * * * * * End of Message * * * * *
5  Sending Query to Name Server at 9.67.1.5 Result: OK
6  Notification Arrived: UDP data delivered RC = OK
7  UDP Data Length: 55
   Return from WaitForAnswer: OK
   * * * * * Beginning of Message * * * * *
   Query Id:                 1
   Flags:                    1000 0101 1000 0000
   Number of Question RRs:   1
   Question   1: RICKA.RALEIGH.IBM.COM A IN
   Number of Answer RRs:     1
8  Answer     1: RICKA.RALEIGH.IBM.COM 86400 A IN 9.67.97.3
   Number of Authority RRs:  0
   Number of Additional RRs: 0
   * * * * * End of Message * * * * *
   HostNumber (1) is: 9.67.97.3
```

Figure 63. Example of RESOLVER trace output

Following are short descriptions of the numbered items in the trace.

1 Address of the name server being used for name resolution. The address is pulled from the TCPIP.DATA data set.

2 Identification number of the query. This is also returned in the response and should be used to match queries to responses.

3 Bits set to determine the type of query and response. (See RFC 1035.) There are 16 bits (0–15) set in the parameter field of DNS message.

Bit	Meaning
0	Operation: 0=query, 1=response
1–4	Query type: 0=standard, 1=inverse
5	Set if the answer is authoritative
6	Set if the message is truncated
7	Set if recursion is wanted
8	Set if recursion is available

9–11 Reserved

12–15 Response type

Value	Meaning
0	No error
1	Format error in query
2	Server failure
3	Name does not exist

4 Actual question sent to the name server

5 IP address of the name server being queried

6 The response has arrived (UDP in this case)

7 Length of the record

8 Answer to the question

Chapter 18. Diagnosing z/OS UNIX sendmail and popper problems

This topic describes how to diagnose problems with z/OS UNIX sendmail, an electronic mail-transport agent and server, and with z/OS UNIX popper, a mail-delivery agent.

The following subtopics are in this topic:

- "Diagnostic aids for sendmail"
- "Debugging switches"
- "Additional diagnostic aids" on page 549
- "Diagnostic aids for IPv6 support" on page 551
- "Diagnostic aids for AT-TLS support" on page 552
- "Diagnostic aids for mail filter support" on page 552
- "Hints and troubleshooting sendmail message submission program (MSP) file submit.cf" on page 553
- "Diagnostic aids for popper" on page 554

Diagnostic aids for sendmail

The following sections describe various tools and techniques available for diagnosing problems with z/OS UNIX sendmail. For a comprehensive discussion of sendmail, see the industry-accepted publication *sendmail* by O'Reilly & Associates, Inc. (ISBN 1-56592-839-3). That publication is known throughout the industry as the *bat book*, because of the fruit bat depicted on the cover. This topic consistently refers to the *bat book* for further information.

You can also find more information about sendmail at the web site.

For information about diagnosing problems with the other z/OS Communications Server mail application, Simple Mail Transfer Protocol (SMTP), see Chapter 17, "Diagnosing Simple Mail Transfer Protocol (SMTP) problems," on page 535.

Debugging switches

Table 40 shows a complete list of debugging switches in sendmail. Some of these switches create long and complex output. Each switch that is especially useful for debugging mail problems is marked "X" in the third column.

Table 40. Debugging switches by category

Category	Bat book reference	Useful for mail problems	Description
-d0.1	16.6.1	X	Print version, compilation, and interface information
-d0.4	16.6.2	X	Our name and aliases
-d0.10	16.6.3		Operating System defines
-d0.12	16.6.4	X	Print library (libsm) defines
-d0.13	16.6.5	X	FFR Defines: _FFR_MILTER_PERDAEMON
-d0.22	16.6.6		Dump delivery agents

© Copyright IBM Corp. 2000, 2015

Table 40. Debugging switches by category (continued)

Category	Bat book reference	Useful for mail problems	Description
-d0.40	16.6.7		Print network address of each interface
-d0.44	16.6.8		End with finis()
-d2.9	16.6.9		Show file descriptors with *dumpfd()*
-d1.1	16.6.10		Trace enoughspace()
-d1.5	16.6.11		Show failed mail
-d2.1	16.6.12		DNS name resolution
-d2.9	16.6.13		Call to getcanonname(3)
-d3.1	16.6.14		Trace dropped local hostnames
-d3.5	16.6.15		Hostname being tried in getcanonname(3)
-d3.15	16.6.16		Yes/no response to -d8.5
-d3.20	16.6.17		Resolver debugging
-d3.30	16.6.18		Trace delivery
-d11.2	16.6.19	X	Show the user-id running as during delivery
-d12.1	16.6.20		Show mapping of relative host
-d13.1	16.6.21		Show delivery
-d20.1	16.6.22		Show resolving delivery agent:parseaddr()
-d21.1	16.6.23	X	Trace rewriting rules
-d21.2	16.6.24		Trace $&macros
-d22.1	16.6.25		Trace tokenizing an address : prescan()
-d22.11	16.6.26		Show address before prescan
-d22.12	16.6.27		Show address after prescan
-d25.1	16.6.28		Trace "sendlist"
-d26.1	16.6.29		Trace recipient queueing
-d27.1	16.6.30	X	Trace aliasing
-d27.2	16.6.31	X	Include file, self-reference, error on home
-d27.3	16.6.32	X	Forwarding path and alias wait
-d27.4	16.6.33	X	Print not safe
-d27.5	16.6.34	X	Trace aliasing with printaddr[]
-d27.8	16.6.35	X	Show setting up an alias map
-d27.9	16.6.36		Show user-id/group-id changes with:Include:reads
-d28.1	16.6.37		Trace user database transactions
-d29.1	16.6.38		Special rewrite of local recipient
-d29.4	16.6.39		Trace fuzzy matching
-d31.2	16.6.40		Trace processing of headers
-d34.1	16.6.41		Watch header assembly for output
-d34.11	16.6.42		Trace header generation and skipping
-d35.9	16.6.43		Macro values defined
-d37.1	16.6.44	X	Trace settings of options
-d37.8	16.6.45	X	Trace adding of words to a class

Table 40. Debugging switches by category (continued)

Category	Bat book reference	Useful for mail problems	Description
-d38.2	16.6.46		Show database map opens and failures
-d38.3	16.6.47	X	Show passes
-d38.4	16.6.48	X	Show result of database map open
-d38.9	16.6.49		Trace database map closing and appends
-d38.10	16.6.50		Trace NIS search for @:@
-d38.12	16.6.51		Trace database map stores
-d38.19	16.6.52		Trace switched map finds
-d38.20	16.6.53		Trace database map lookups
-d41.1	16.6.54		Trace queue ordering
-d44.4	16.6.55		Trace safefile()
-d44.5	16.6.56		Trace writable()
-d48.2	16.6.57		Trace calls to the check_rules set
-d49.1	16.6.58		Trace checkcompat()
-d52.1	16.6.59		Show disconnect from controlling TTY
-d52.100	16.6.60		Prevent disconnect from controlling TTY
-d60.1	16.6.61		Trace database map lookups inside rewrite()
-d99.100	16.6.62		Prevent backgrounding including the daemon
-d96.9	NA	X	Trace SSL (gsk_xxx) calls

Additional diagnostic aids

In addition to debugging switches, you can use the following z/OS UNIX sendmail diagnostic aids:

- syslog.log provides more information. The following sample shows a z/OS UNIX sendmail syslog.log message:

  ```
  Dec 28 02:13:30 MVS186 sendmail[67108947]: EZZ7514I: sendmail starting
  .
  .
  Dec 28 02:13:30 MVS186 sendmail[67108947]: starting daemon (8.12.1): SMTP
  ```

  For descriptions of sendmail messages, see z/OS Communications Server: IP Messages Volume 4 (EZZ, SNM).
- Use the -v (verbose) command-line switch to print a complete description of all the steps required to deliver a mail message. For details, see *sendmail, 3rd Edition*.
- Use the -X (trace log) command-line switch to record all input, output, SMTP traffic, and other significant transactions into the specified trace file. For details, see *sendmail, 3rd Edition*.
- Check the qf file for queueing concerns. z/OS UNIX sendmail stores undeliverable messages in the QueueDirectory that is specified in the configuration file. The QueueDirectory contains data files (df files) named `dfxxxxxxxx` and matching queue-control files (qf files) named `qfxxxxxxxx`. A df file contains the body of a queued message. A qf file holds all the information that is needed to deliver the message. Each queued message has a corresponding df and qf file.

The qf file is line-oriented, containing one item of information per line. The single uppercase character (the code letter) specifies the contents of the line. The complete list of qf code letters is shown in Table 41.

Table 41. qf File code letters

Code	Reference	Meaning	How Many
A	Bat book 11.11.1	AUTH=parameter	At most, one
B	Bat book 11.11.2	Message body type	At most, one
C	Bat book 11.11.3	Set controlling user	At most, one per R line
d	Bat book 11.11.4	Data file directory	Exactly one
D	Bat book 11.11.5	Data file name	Exactly one
E	Bat book 11.11.6	Send errors to	Many
F	Bat book 11.11.7	Save flagged bits	Exactly one
H	Bat book 11.11.8	Header line	Many
I	Bat book 11.11.9	Mode and device information for the df file	Exactly one
K	Bat book 11.1.10	Time last processed	Exactly one
M	Bat book 11.11.11	Message (why Manyqueued)	At most one
N	Bat book 11.11.12	Number times tried	At most, one
P	Bat book 11.11.13	Priority (current)	At most, one
Q	Bat book 11.11.14	The DSN ORCPT address	At most, one per R ine
r	Bat book 11.11.15	Final recipient	At most, one
R	Bat book 11.11.16	Recipient address	Many
S	Bat book 11.11.17	Sender address	Exactly one
T	Bat book 11.11.18	Time created	Exactly one
V	Bat book 11.11.19	Version	Exactly one
Z	Bat book 11.11.20	DSN envelope ID	At most, one
!	Bat book 11.11.21	Delivery by specification	At most, one
$	Bat book 11.11.22	Restore macro value	At most, one

Table 41. qf File code letters (continued)

Code	Reference	Meaning	How Many
.	Bat book 11.11.23	End of qf file	Exactly one

Bat book refers to *sendmail* by O'Reilly & Associates, Inc. (1-56592-839-3). Op refers to *Sendmail Installation and Operation Guide* that is shipped in /usr/lpp/tcpip/samples/sendmail/sendmail.ps.

Diagnostic aids for IPv6 support

For information about configuring an IPv6 Daemon, see z/OS Communications Server: IP Configuration Guide.

In addition, to handle network variation, the following are useful.

- Failed to open socket.

  When invoking sendmail, if it fails to open a socket, the following log message is displayed:

  ```
  opendaemonsocket: daemon < MTA_name>: cannot create server SMTP socket"
  opendaemonsocket: daemon <MTA_name>: problem creating SMTP socket"
  ```

  Consider the following to solve this problem:

  – Is the TCP/IP stack enabled for IPv4 or IPv6?
  – Is the DaemonPortOption in sendmail configuration file (sendmail.cf) properly set? (Remember that an IPv6 daemon option cannot run on a IPv4-only stack.)

- DNS support.

  When sendmail runs as a IPv6-enable daemon, it needs to do two things:

  – Receive mails with long-type address
  – Make AAAA type queries with DNS

  In some database files (for example, aliases, relay-domains, or access), if mail which is targeted to a legal IPv6 site always fails to be sent, check whether name server supports IPv6 (AAAA type queries).

  If DNS queries are failing, see "RESOLVER trace (SYSTCPRE)" on page 170 for information about how to run a resolver trace.

  To determine whether the name server is IPv6-capable, issue the following:

  ```
  dig @<address_of_name_server> <host_name_of_target> aaaa
  ```

  If this does not return an IPv6 address, either the name server is not IPv6-capable or the name server is not configured properly.

  To determine whether the name server is IPv6-capable and the name server is a bind-based name server (not Microsoft), issue the following:

  ```
  dig @<address_of_name_server> version.bind chaos any
  ```

  If the version of bind returned is 9.0 or greater, the name server is IPv6-capable, so it is likely not configured properly for IPv6. DNS administrators can restrict the name server from giving out its bind version, but if any type of an answer is received other than a failed query response, the name server is IPv6-capable. If the query fails, the name server cannot support IPv6.

Diagnostic aids for AT-TLS support

Before you begin, you need to know that a packet trace can be taken to ensure that mail is encrypted before being sent. If packet traces show that encryption has occurred, but a specific packet is suspected of being unencrypted, set confLOG_LEVEL to a value greater than 9 and re-create the packet. If there were any errors in encryption, they are sent to syslog with LOG_ERR. After investigating a single packet, if you want to investigate whether SSL function calls were in error, use -d96.9 debug to check all return codes to gsk_xxx calls.

To analyze the reason individual System SSL function calls are in error, follow these steps:

1. Set the /etc/mail/zOS.cf file GskTraceFile parameter to a file name to receive the System SSL trace.
2. Rerun the command.
3. Use the System SSL **gsktrace** command to create a readable copy of the trace information.

When you are done, you can use this trace information to analyze reasons individual System SSL function calls might be in error. For additional information, see z/OS Cryptographic Services System SSL Programming.

Diagnostic aids for mail filter support

The debug message of a mail filter can be divided into two parts:

- Milter API

  These messages are provided to allow programmers to develop a mail filter. These messages are written into the log file defined in filter program. The following section gives more detail of these messages.
- Filter program

  The Milter API messages are mainly function error and input error. A function error means that a function call fails. It occurs when using an incompatible function or allocating invalid system resource, for example. These messages can be as follows:

```
                :
EZZ9963I filtername: malloc(size) failed for typestorage (ret reason)
   strerror(ret) {abort | try again}"
EZZ9971I filtername: pthread_create() failed (ret reason), strerror(ret)
                :
```

These errors cannot be resolved easily. Report them to the program developer or the system administrator.

An input error means that a user has given an invalid parameter and caused the program to terminate. The mail filter reads socket type and port number from users.

Socket type has the following types:

- inet4 (for IPv4)
- inet6 (for IPv6)
- UNIX domain socket

The following list describes the error operation and messages:

EZZ9951I SampleFilter: unknown socket type inet5
You gave an invalid socket type inet5. Select a valid socket type.

EZZ9961I filtername: Unable to bind (ret reason) to port stringI: strerror(ret)
The file path does not exist when using UNIX domain socket. Check that the file path exists before using UNIX domain socket.

EZZ9952I filtername: UNIX socket name string longer than max
A UNIX domain socket name cannot be defined over 108 characters in length. Rename A.B to less than 108 characters.

EZZ9955I SampleFilter: unknown port name abc
You gave an invalid port number.

EZZ9961 filtername: Unable to bind (ret reason) to port string: strerror(ret)
Do not give a port number that has been reserved, for example 21 (default for FTP). Obtain the reserved filter port number from the system administrator.

EZZ9965I SampleFilter: Unable to create listening socket on conn inet:21
Some error occurred when creating, binding, setting or listening a socket. Detailed error message should already be displayed before this message.

Sendmail daemon provides some information for connecting and talking to mail filters, you can change the log level defined in sendmail configuration file. The default log level is the same with sendmail log level:

```
O Milter.LogLevel=20
```

Check if the sendmail daemon works correctly with mail filters by log messages in sendmails's log file.

O Milter.LogLevel=20
Check if sendmail daemon works correctly with mail filters in the sendmail log file.

If mail is lost between the sendmail daemon and the filter program, see "Packet trace (SYSTCPDA) for TCP/IP stacks" on page 93 to run a packet trace to determine where, and if, packets are being lost.

Hints and troubleshooting sendmail message submission program (MSP) file submit.cf

When feature msp is specified, FEATURE('msp'), the option conf RunAsUser is set to smmsp. This user must have the group smmspgrp, for example, the same group as the clientmqueue directory. If you specify a user whose primary group is not the same as that of the clientmqueue directory, then you should explicitly set the group as follows:

```
FEATURE('msp')
define('confRUN_AS_USER', 'mailmsp:smmspgrp')
```

The SEZASAMP(EZARACF) file shows sample commands to add the smmsp user and group.

```
ADDGROUP SMMSPGRP OMVS(GID(25))
ADDGROUP SNDMGRP OMVS(GID(26))
ADDUSER MAILNULL DFLTGRP(SNDMGRP) NOPASSWORD OMVS(UID(26) HOME('/'))
ADDUSER SENDMAIL DFLTGRP(SNDMGRP) NOPASSWORD OMVS(UID(0) HOME('/'))
ADDUSER SMMSP DFLTGRP(SMMSPGRP) NOPASSWORD OMVS(UID(25) HOME('/'))
```

In addition, there are security concerns for programs that change user ID without prompting for a password. Program control is the Security Server facility used to

manage programs that change user IDs without prompting for a password. By having an installation use program control, applications not permitted to the facility are not allowed to change user IDs without prompting for a password. The commands are:

```
PERMIT BPX.DAEMON CLASS(FACILITY) ID(SENDMAIL) ACCESS(READ)
SETROPTS RACLIST(FACILITY) REFRESH
```

For more information about Security Server commands used to allow sendmail access to the program control facility, see SEZAINST(EZARACF). For complete information about the program control facility, see z/OS Security Server RACF Security Administrator's Guide.

In a program control environment, use /bin/sendmail to create mail as a Mail Submission Agent (MSA) and /usr/sbin/sendmail as a Mail Transfer Agent (MTA).

If a program control environment is defined for your installation and an end user invokes sendmail and gets EZZ9895I, the installation has not configured the MSA completely.

/bin/sendmail must be owned by the same user ID as the confRUN_AS_USER (smmsp uid 25 default) set in /etc/mail/submit.cf. To do this enter the following two commands:

```
chown 25:25 /bin/sendmail
chmod 6755 /bin/sendmail
```

Diagnostic aids for popper

Diagnostic aids for popper are found in the SYSLOGD log information. Following is a sample z/OS UNIX popper log message:

```
Apr 20 14:19:36 MVSW popper[16777240]: Received: "quit"
```

Use the -t trace option to direct all popper message logging to the specified file. The POP server copies the user's entire maildrop to /tmp and then operates on that copy. If the maildrop is particularly large, or inadequate space is available in /tmp, then the server refuses to continue and terminate the connection.

To test popper, you can mimic a popper client by TELNETing into a popper port (110) and issuing the **popper** commands that are documented in RFC 1725. The following are a few of the commands that are used to verify that popper is listening on port 110:

user *name*
: Specifies the mailbox.

pass *string*
: Specifies a server/mailbox-specific password.

list [*msg*]
: Lists all message numbers and size or information about a specific message.

retr *msg*
: Retrieves the specific message to the screen.

quit
: Closes the connection to popper.

Following is an example of a TELNET exchange:

```
> telnet <host name/ip addr> 110
OK POP (version 2.53) at MVSW.tcp.raleigh.ibm.com starting.

> user user163
OK Password required for USER163

> pass tcpxyz
OK USER163 has 6 messages (4273 octets)

> list
OK 6 messages (4273 octets)
1 346
2 371
3 333
4 347
5 2541
6 335
.

> retr 3
OK 333 octets
Received: 9BPXROOT@local host by mvsw.tcp.raleigh.ibm.com (8.8.7/8.8.1) id
PAA83
886099 for user163; Tue, 10 Mar 1998 15:36:57 -0500
Date: Tue, 10 Mar 1998 15:36:57 -0500
from USER163 <USER163>USER163
Message-ID: <199803102036.PAA83886099@mvsw.tcp.raleigh.ibm.com>
X-UIDL: 4569e8e12631e857eed8d8b0ca493
Status: 0

hello

.
```

Chapter 19. Diagnosing SNALINK LU0 problems

The TCP/IP host is implemented with the SNALINK LU0 function. This function allows the use of an SNA backbone to transfer TCP/IP protocols. A TCP/IP host with SNALINK LU0 can be an originator, destination, or router for TCP/IP data. To use the SNALINK LU0 function of TCP/IP, each connected host must have VTAM and TCP/IP installed. The SNALINK LU0 application runs in its own address space and is defined as a VTAM application. There are two types of SNALINK implementations:

- SNALINK LU0, which uses VTAM LU0 protocol
- SNALINK LU6.2, which uses VTAM LU6.2 protocol

This topic describes how to diagnose problems with the SNALINK LU0 function and contains the following sections:

- "Definitions"
- "Problem diagnosis"
- "Traces" on page 14

SNALINK LU6.2 diagnosis is discussed in Chapter 20, "Diagnosing SNALINK LU6.2 problems," on page 565.

SNALINK LU0 is a very convenient way to connect to TCP/IP hosts using an existing SNA backbone. An IP datagram destined for a remote host that is connected using SNALINK LU0 is passed to the SNALINK LU0 address space by TCP/IP. The data is packaged into an SDLC frame and transmitted to the remote host using SNA LU0 protocol. Two SNALINK LU0 applications can be configured to connect using a single, bidirectional session or with two separate sessions (one dedicated to send data in each direction).

Definitions

The following are required to define a SNALINK LU0:

- Device and link definitions in the TCPIP profile
- Home address and routing information
- VTAM application definitions
- Parameters on the PROC used to start SNALINK LU0

For more information about these required definitions, see z/OS Communications Server: IP Configuration Reference.

Problem diagnosis

SNALINK LU0 problems are normally reported as one of the following:

- Abends
- Session hung terminals
- Session outages

Use the information in the following sections for problem determination and diagnosis of errors reported against SNALINK LU0.

© Copyright IBM Corp. 2000, 2015

When contacting the IBM Software Support Center for any type of SNALINK LU0 problem, have the VTAM application definitions for SNALINK LU0 and the DEVICE and LINK information from the *hlq*.PROFILE.TCPIP data set for SNALINK LU0.

Abends

An abend for the SNALINK LU0 application should result in messages or error-related information about the MVS system console. Since SNALINK LU0 is a VTAM application, some abends might be generated or first detected by VTAM. These messages indicate that VTAM is abending or a dump is being taken for the SNALINK LU0 application.

If a VTAM error caused by SNALINK LU0, see z/OS Communications Server: SNA Messages and z/OS Communications Server: IP and SNA Codes for initial problem determination.

If SNALINK LU0 fails to initialize with an 0C4 abend, there is probably an installation problem. Check the program properties table (PPT) entries for errors. Some levels of MVS do not flag PPT syntax errors properly. For more information about PPT configuration, see z/OS MVS Initialization and Tuning Reference.

Documentation

Code a SYSMDUMP DD or SYSABEND DD statement in the SNALINK cataloged procedure.

There are two MVS abends commonly seen during the initialization and startup of the SNALINK LU0 application: X'0C2' and X'0F8'. Both can be caused by the SNALINK LU0 application processing in TCB mode. The VTAM application definition statement for SNALINK LU0 must have the SRBEXIT=YES parameter coded. This should ensure that VTAM passes control to SNALINK LU0 in SRB mode. SNALINK LU0 code has processing that is not allowed in TCB mode. If the SRBEXIT parameter is coded incorrectly or allowed to default, either abend X'0C2' or X'0F8' will occur.

Guideline: Some networking optimizing packages change the defined mode for VTAM applications for performance purposes. It is suggested that this type of program not be used for the SNALINK LU0 application.

Analysis

For more information about debugging abends, see Chapter 3, “Diagnosing abends, loops, and hangs,” on page 21.

An abend or unexpected termination of the SNALINK LU0 application does not terminate the TCP/IP address space. If there is no alternate route to the remote host, IP datagrams for TCP/IP Services components (such as TELNET and FTP) are not transmitted until the application is restarted, either manually or using TCP/IP autolog.

Session hangs

This section discusses diagnosis of a hung terminal after a session has been successfully connected. A hang might be detected by TCP/IP users who are connected to the remote system with SNALINK LU0 (this could be FTP, TELNET, or other applications).

The SNALINK LU0 application detects a hung terminal if there is no response to data sent. After waiting 30 seconds for a response, SNALINK LU0 ends the session and tries to reestablish the LU-to-LU session with its partner SNALINK LU0 application. This processing is shown on the SNALINK LU0 log or MVS console log.

Documentation

To determine the cause of an SNALINK LU0 session hung terminal, the following might be needed:

- SNALINK LU0 log or MVS console log
- NETSTAT DEVLINKS display output
- VTAM display application status output
- SNALINK LU0 DEBUG trace output
- VTAM buffer trace of the SNALINK LU0 applications
- VTAM internal trace

For information about VTAM traces, see 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.

This list of documentation includes documentation needed to resolve most types of hung terminals. All of the indicated data might not be needed for each occurrence of a hung terminal. The following section provides information about the types of data that might be needed for each diagnostic step.

Steps for analyzing session hangs

Determine the cause of a reported SNALINK LU0 hung terminal.

Before you begin

The first step in analysis is to determine whether the SNALINK LU0 is actually hung or if one of the sessions using SNALINK LU0 to transfer data is hung. When the SNALINK LU0 is the only connection between two hosts, an actual hang in the SNALINK LU0 application impacts all data flowing for TCP/IP. This can include TELNET, FTP, and any other application.

Procedure

Perform the following steps:

1. Does all traffic across the SNALINK LU0 stop? A VTAM buffer trace of the SNALINK LU0 application can be used to see whether any data is being passed. If data is still flowing on the session, the SNALINK LU0 is not hung. You need to determine which TCP/IP application or component is failing. If there is no data traffic, continue with Step 2. You can also check SNALINK LU0 traffic by doing multiple VTAM displays of the SNALINK LU0 application. The SEND and RECEIVE data count should increase for an active session. Often, using the VTAM display to obtain the status of the TRLE might provide useful information.
2. Issue NETSTAT DEVLINKS to determine the status of the SNALINK LU0 TCP/IP device. If the NETSTAT output shows that the application is trying to connect, check the VTAM and SNALINK LU0 consoles for information about a previous error or abend. If NETSTAT indicates "negotiating," verify the session type. You might require a session_type of SINGLE; see the z/OS

Communications Server: IP Configuration Reference for information about configuring session types. If NETSTAT indicates “connected” or “sending,” continue with Step 3.

3. At this point, you should determine the last SNALINK LU0 activity or processing. This is best accomplished with the debug trace. Contact your IBM Software Support Center with information about the last activity from the SNALINK LU0 console and debug trace.

Session outages

A session outage is an unexpected abend or termination of the task. Session outages are seen only when an unrecoverable error is detected. The error might be a SNALINK LU0 abend or an error return code from a VTAM request. A session outage should not occur without an indication of its cause, either on the SNALINK LU0 or the VTAM console. Since SNALINK LU0 abends were already described separately, this topic describes other types of session outages.

For an example of a successful session setup between two SNALINK LU0 applications, see the z/OS Communications Server: IP Configuration Reference.

Documentation

The following documentation might be needed to determine the source of the error for a session outage problem:

- SNALINK LU0 log
- MVS console log
- VTAM log
- NETSTAT DEVLINKS display output
- VTAM display application status output
- SNALINK LU0 DEBUG trace output
- VTAM buffer trace of the SNALINK
- LU0 applications
- VTAM Internal Trace (VIT)

Note: For information about VTAM traces, see 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.

Analysis

When a SNALINK LU0 outage occurs, there should be messages and indicators of the reason for the outage. These appear in the SNALINK LU0 log, or on the VTAM console, or both. If an abend has been recorded, continue diagnosis using the section on abends.

The following is an example of a session outage problem. The message EZA5797E `Rejecting bind from xxxxx-no DLC found`, along with VTAM error message IST663I `Bind fail request received, SENSE=080A0000`, was displayed on the MVS system console.

Cause: Large packet size sent in a PIU is rejected by the NCP with sense 800A0000 (PIU too long).

Resolution: Reduce the MTU size on this route by using the BEGINROUTES statement.

Traces

The following traces are useful:

- Use VTAM buffer trace to trace the data sent and received from the VTAM.
- Use the TCPIP PKTTRACE LINKNAME=link_name to trace the data sent and received from TCP/IP.

Using IP packet trace

The IP packet trace facility is used to trace the flow of IP packets. It is useful when tracking the cause of packet loss or corruption. If the LINKNAME parameter of the IP packet trace facility is specified, only packets that are transferred along the specific link are traced. Specifying this parameter is recommended to avoid tracing many unrelated packets. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for details about how to use the IP packet trace facility.

SNALINK LU0 DEBUG trace

The SNALINK LU0 DEBUG trace output is written to an internal buffer. The trace can be seen only if a dump of the SNALINK LU0 address space is taken. The trace wraps when the buffer is full (a pointer in the trace header points to the most current entry).

The trace contains information about SNALINK LU0 processing. This information includes communication with VTAM and TCP/IP, showing VTAM macro requests and DLC requests.

Starting SNALINK LU0 DEBUG trace

To run the SNALINK LU0 DEBUG trace, SNALINK LU0 must be started with DEBUG listed as the first parameter of the PARM parameter on the EXEC statement of the SNALINK cataloged procedure. For information about this parameter, see z/OS Communications Server: IP Configuration Reference.

DEBUG trace example

Figure 64 on page 562 shows part of an internal SNALINK LU0 trace obtained from a dump. As shown in the example, the trace can be located by searching for the characters TRCTBL in the dump of the SNALINK LU0 address space. Following the eyecatcher is the address of the next entry to be written, the starting address of the trace table, and the ending address of the trace table.

Use the information following the trace to interpret the entry types and their meaning.

```
00012A40   00000000   00000000   E3D9C3E3   C2D34040   | ........TRCTBL   |
00012A50   00018640   00018300   00020000   A020006C   | ..f ..c........% |
00012A60   00000000   80C15008   94000001   00000000   | .....A&.m....... |
.
.
.
00018300   B56D355D   F3C92348   00000000   00000000   | ._.)3I.......... |
00018310   40000000   00000045   00000000   800091DA   |  ............j.  |
00018320   B56D355D   F76940CA   00000000   00000000   | ._.)7. ......... |
00018330   10230000   00012B80   00000000   00000000   | ................ |
00018340   B56D355D   F769568A   00000000   00000000   | ._.)7........... |
00018350   40000000   00000006   00000002   8000E256   |  .............S. |
00018360   B56D355D   F77C748A   00000000   00000000   | ._.)7@.......... |
00018370   40000000   00000007   00000001   8000E2A2   |  .............Ss |
00018380   B56D355D   F789628A   00000000   00000000   | ._.)7i.......... |
00018390   40000000   0000005A   00000002   8000E2D4   |  ......!......SM |
000183A0   B56D356C   173D7DC8   E2D5C1D3   E4F0F2C1   | ._.%..'HSNALU02A |
000183B0   40000000   00000034   00000000   8000CEF8   |  .............8  |
000183C0   B56D356C   1788D188   E2D5C1D3   E4F0F2C1   | ._.%.hJhSNALU02A |
000183D0   10170000   00024E30   00000000   80000000   | ......+......... |
000183E0   B56D356C   58C0DACB   00000000   00000000   | ._.%.{.......... |
000183F0   17101001   00024E30   00024FC0   080A0000   | ......+...|{.... |
00018400   B56D3573   3F9CF1CB   00000000   00000000   | ._....1......... |
00018410   31000800   091BF1F8   31010207   00000000   | ......18........ |
00018420   B56D3573   3F9E5C4B   E2D5C1D3   E4F0F2C1   | ._....*.SNALU02A |
00018430   40000000   00000019   00000000   80009E74   |  ..............  |
00018440   B56D3573   3FDE2DCB   00000000   00000000   | ._.............. |
00018450   2A100000   00024DC0   00024F80   00000000   | ......({..|..... |
00018460   B56D3573   3FDEF7CB   E2D5C1D3   E4F0F2C1   | ._....7.SNALU02A |
00018470   102A0000   00024DC0   00024F80   80000000   | ......({..|..... |
00018480   B56D3573   3FDF0C8B   E2D5C1D3   E4F0F2C1   | ._......SNALU02A |
00018490   40000000   00000034   00000000   8000CEF8   |  .............8  |
000184A0   B56D3573   401997CB   E2D5C1D3   E4F0F2C1   | ._.. .p.SNALU02A |
000184B0   10170000   00024E30   00024FC0   80000000   | ......+...|{.... |
000184C0   B56D3573   4019F1CB   E2D5C1D3   E4F0F2C1   | ._.. .1.SNALU02A |
000184D0   40000000   00000024   00000000   8000A59C   |  .............v. |
000184E0   B56D3573   8161A8CB   00000000   00000000   | ._..a/y......... |
000184F0   17100000   00024E30   00024FC0   00000000   | ......+...|{.... |
.
.
.
000185F0   23100000   00024F10   00037000   80000118   | ......|......... |
00018600   B56D3585   4906F5C0   E2D5C1D3   E4F0F2C1   | ._.e..5{SNALU02A |
00018610   02000000   03790011   00037000   84000C60   | .....`......d..- |
00018620   B56D3585   49081240   E2D5C1D3   E4F0F2C1   | ._.e... SNALU02A |
00018630   10230000   00024F10   00037000   80009000   | ......|......... |
00018640.:018FFF.--All bytes contain X'00'
```

Figure 64. Example of a SNALINK LU0 DEBUG trace

The layout of a SNALINK trace table entry is shown in Table 42.

Table 42. Format of a SNALINK trace table entry

Bytes	Definition
00–07	TOD time stamp
08–0F	LU name, if any

Table 42. Format of a SNALINK trace table entry (continued)

Bytes	Definition
10	Entry Type **Value** 01 DLC Accept 02 DLC Send 03 DLC Receive 04 DLC Sever 05 DLC Msg Pend Queue Request 06 DLC Msg Pend D-Queue Request 0E MVS DLC emulation 0F DLC Interrupt 10 VTAM Request 17 VTAM OPNDST Exit 1F VTAM CLSDST Exit 22 VTAM SEND Exit 23 VTAM Receive Exit 25 VTAM SESSIONC Exit 2A VTAM OPNSEC Exit 2C VTAM TERMSESS Exit 31 VTAM SCIP Exit 32 VTAM LOSTERM Exit 33 VTAM NSEXIT Exit 34 VTAM TPEND Exit 35 VTAM LOGON Exit 40 SNALINK Internal Message Routine Call
11	DLC Interrupt Code/VTAM RPL REQ Code/ VTAM Receive Exit Chain field
12	VTAM CMD: R15/VTAM Exit: RTNCD
13	VTAM CMD: R0 /VTAM Exit: FDB2/DLC IPRCODE
14-17	RPL Address/DLC MSG ID/TPEND reason code/Internal Message ID
18–1B	VTAM Send/Receive/DLC buffer address/Number of Arguments Passed to Internal Message routine
1C–1F	VTAM Send/Receive/DLC buffer length/Internal Message Routine caller's return address

Chapter 20. Diagnosing SNALINK LU6.2 problems

This topic describes how to diagnose problems with the SNALINK LU6.2 function and contains the following subtopics:

- "Steps for setting up a SNALINK LU6.2 network" on page 566
- "Common configuration mistakes" on page 567
- "Diagnosing problems" on page 568
- "Documentation references for problem diagnosis" on page 579
- "Traces" on page 584
- "Finding abend and sense code documentation" on page 589
- "Finding error message documentation" on page 589

The SNALINK LU6.2 interface uses the LU type 6.2 protocol to establish a point-to-point connection across a SNA network. SNALINK LU6.2 is capable of establishing a connection with any system that runs TCP/IP and uses the LU type 6.2 protocol.

The SNALINK LU6.2 interface is similar to the SNALINK LU0 and X.25 NPSI interfaces with the connection involving several subsystems. The components of the SNALINK LU6.2 network are shown in Figure 65.

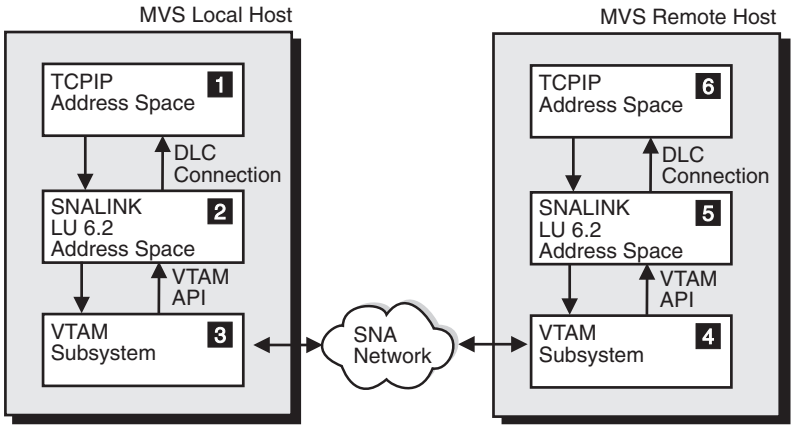

Figure 65. Components of a SNALINK LU6.2 connection on MVS

Following is a brief description of the component interaction and data flow that occurs when data is transferred over a SNALINK LU6.2 network. Each component is cross-referenced to the figure.

1 Data is generated and encapsulated on the TCP/IP address space and is passed to the SNALINK LU6.2 address space through a DLC connection.

2 The SNALINK LU6.2 address space handles all establishment, aging, and termination of SNA network connections in a manner transparent to the TCP/IP address space. The data is then sent to the local system SNA subsystem. f MVS hosts, this subsystem is VTAM.

© Copyright IBM Corp. 2000, 2015

3 VTAM APPC routines are used to pass the data to the SNA network.

4 VTAM routines on the destination system receive the data and pass it through to the SNALINK LU6.2 address space.

5 The SNALINK LU6.2 address space sends the data to the TCP/IP address space using a DLC connection.

6 The data is unencapsulated and processed by the TCP/IP address space.

Steps for setting up a SNALINK LU6.2 network

This list of steps can be used to diagnose problems in starting components by identifying the prerequisites. For details about how to complete the steps, see the appropriate documentation.

Procedure

Complete the following steps to establish the system that is described in Figure 65 on page 565.

1. Configure the SNALINK LU6.2 network on both the local and remote network hosts. This is fully described in z/OS Communications Server: IP Configuration Reference in the section about configuring and operating the SNALINK LU6.2 interface. The process can be condensed into the following steps:
   a. Specify SNALINK LU6.2 DEVICE and LINK statements in the *hlq*.PROFILE.TCPIP data set.
   b. Copy the sample SNALINK LU6.2 cataloged procedure to an authorized data set and update according to your system.
   c. Define a SNALINK LU6.2 application LU to VTAM.
   d. Customize a SNALINK LU6.2 configuration data set.
2. Vary the SNALINK LU6.2 VTAM application LUs active on both the local and remote network hosts.
3. Start both the local and remote network TCP/IP address spaces.
4. Start both the local and remote network SNALINK LU6.2 address spaces, if they have not been autologged by the TCP/IP address space.
5. Verify that the network connection has been established between the local host and the remote host. See "Using the SNALINK LU6.2 subcommand" on page 580 for details about how to verify SNALINK LU6.2 connections.

Results

The example in Figure 66 on page 567 shows the messages that are expected when the SNALINK LU6.2 address space is started and a network connection is established.

```
S SNAL621A
$HASP100 SNAL621A ON STCINRDR
$HASP373 SNAL621A STARTED
 1 IEF403I SNAL621A - STARTED - TIME=15.26.03
 2 EZA5927I LU62CFG : NO ERRORS DETECTED - INITIALIZATION WILL CONTINUE
 3 EZA5932I INITIALIZATION COMPLETE - APPLID: SNAL621A  TCP/IP: TCPCS
 4 EZA5935I SEND CONVERSATION ALLOCATED FOR 9.67.22.2
 5 EZA5933I LINK SNALU62L OPENED
EZZ4313I INITIALIZATION COMPLETE FOR DEVICE SNALU621
 4 EZA5936I RECEIVE CONVERSATION ALLOCATED FOR 9.67.22.2
```

Figure 66. Sample MVS System Console Messages on SNALINK LU6.2 Address Space Startup

The following list explains the MVS system console messages on SNALINK LU6.2 address space startup as shown in Figure 66.

1 The SNAL621A address space has been started.

2 The SNALINK LU6.2 configuration data set for the SNAL621A address space has been successfully parsed.

3 The SNAL621A address space displays its local VTAM application LU and the TCP/IP address space name to which it connects.

4 The SNAL621A address space establishes a network connection through the VTAM API.

5 The SNAL621A address space establishes a DLC connection with its TCP/IP address space.

Common configuration mistakes

Following is a list of common configuration mistakes:

- The SNALINK LU6.2 configuration data set contains a syntax error.
- The SYSTCPD or LU62CFG ddnames in the SNALINK LU6.2 cataloged procedure have been assigned to a data set that is not valid.
- The SNALINK LU6.2 VTAM application LU has not been activated.
- The SNALINK LU6.2 VTAM application LU definition has the option SRBEXIT=YES.
- The SNALINK LU6.2 VTAM application LU definition does not have the option APPC=YES.
- The SNALINK LU6.2 VTAM application LU definition specifies a logon mode table in the MODETAB parameter that does not contain the log mode entry specified in the LOGMODE parameter on the LINK statement in the SNALINK LU6.2 configuration data set. The logon mode entry options used for the local host must be the same as for the remote host.
- The *hlq*.PROFILE.TCPIP data set contains syntax errors in the SNALINK LU6.2 BEGINROUTES, DEVICE, LINK, HOME, or START statements.
- The maximum buffer size in the SNALINK LU6.2 configuration data set does not match the MTU size on the BEGINROUTES statement of the *hlq*.PROFILE.TCPIP data set.
- The link name in the SNALINK LU6.2 configuration data set does not match the link name on the LINK statement in the *hlq*.PROFILE.TCPIP data set.
- The SNALINK LU6.2 device has not been started by a START statement in the *hlq*.PROFILE.TCPIP data set.
- The user ID assigned to the SNALINK LU6.2 start procedure has not had an OMVS Segment assigned to it using RACF or similar security manager.

Diagnosing problems

SNALINK LU6.2 problems are normally reported under one of the following categories:

- Problems starting the SNALINK LU6.2 address space
- DLC connection
- Network connection establishment
- Network connection loss
- Data loss
- Data corruption

Use the information in the following topics to help you diagnose SNALINK LU6.2 problems.

Quick checklist for common problems

The following list summarizes some initial checks that can be made quickly.

Use the following checklist to identify problem areas:

1. **Is the TCP/IP SNALINK LU6.2 network active?**

   PING the remote TCP/IP host from the local TCP/IP host to verify that the SNALINK LU6.2 network is active. If the SNALINK LU6.2 network is not active, continue through this list to identify the problem.

   If the PING still fails after working through this list, see "Network connection establishment problems" on page 574 for a detailed list of network connection problems and their solutions.

2. **Have you completed all the required definitions?**

   See "Steps for setting up a SNALINK LU6.2 network" on page 566 for the list of definitions and configurations required. Continue through this list if connection problems persist.

3. **Have the VTAM major node and application LU used by the SNALINK LU6.2 address space been varied active?**

   See "Useful VTAM operations" on page 581 for details on how to use the VTAM DISPLAY command to identify the status of the VTAM major node and application LU.

   If the VTAM application LU is not in a CONCT state, see "Useful VTAM operations" on page 581 for details about how to vary the VTAM application LU active.

4. **Are the TCP/IP and SNALINK LU6.2 devices started and active on the local and remote host?**

   Check to see whether the TCP/IP and SNALINK LU6.2 devices are active and running. The MVS SDSF facility can be used to view the active address space list for MVS hosts.

   If the SNALINK LU6.2 address space does not start, see "Problems starting the SNALINK LU6.2 address space" on page 569 for a detailed list of startup problems and their solutions.

5. **Did the SNALINK LU6.2 address space list any configuration errors to the SYSPRINT data set?**

   Use the JCL DD statement in the SNALINK LU6.2 cataloged procedure to identify the destination of the SYSPRINT output and check for errors. If errors occur, see "Finding error message documentation" on page 589 to determine

the reason for the configuration errors. Text in the message documentation specifies the action required to fix the problem.

6. **Have the TCP/IP-to-SNALINK LU6.2 DLC connections been established?**

   See "Using NETSTAT" on page 580 for details about how to use the NETSTAT command to identify the status of the DLC connection.

   If the status of the DLC connection is not "Connected," see "DLC connection problems" on page 571 for a detailed list of SNALINK LU6.2 DLC connection problems and their solutions.

7. **Does the MVS system console contain VTAM error messages?**

   See z/OS Communications Server: SNA Messages and z/OS Communications Server: IP and SNA Codes for detailed descriptions of the VTAM error messages and sense codes. These messages might indicate a network configuration or hardware error.

Problems starting the SNALINK LU6.2 address space

Generally, if there is a startup problem, error messages are displayed on the MVS system console during the starting of the SNALINK LU6.2 address space. The address space then terminates.

Documentation

To isolate a SNALINK LU6.2 address space starting problem, note any error messages or abend codes that are displayed on the MVS system console.

Analysis

Table 43 shows some of the common SNALINK LU6.2 address space startup problems.

Table 43. Common SNALINK LU6.2 address space startup problems

If this is displayed. . .	Then this might have occurred . . .	Resolution
The message `Errors Detected - Address Space will Terminate` has been displayed on the MVS system console with no other error messages.	This error message indicates that an error has occurred with the SNALINK LU6.2 configuration data set	Check the SNALINK LU6.2 SYSPRINT output for messages that tell what kind of syntax error might have occurred. If a syntax error has occurred in the configuration data set, correct it and restart the SNALINK LU6.2 address space. See the z/OS Communications Server: IP Configuration Reference for details about the SNALINK LU6.2 configuration data set statement syntax.
The message `Error in open of LU62CFG - no data will be read` has been displayed on the MVS system console.	The SNALINK LU6.2 address space cannot access a SNALINK LU6.2 configuration data set. The LU62CFG ddname might have been omitted from the SNALINK LU6.2 cataloged procedure.	Check the SNALINK LU6.2 cataloged procedure. Ensure that the LU62CFG ddname is assigned a valid SNALINK LU6.2 configuration data set. See the z/OS Communications Server: IP Configuration Reference for an example of a SNALINK LU6.2 cataloged procedure.

Table 43. Common SNALINK LU6.2 address space startup problems (continued)

If this is displayed. . .	Then this might have occurred . . .	Resolution
The message `Address Space Already Active - this Address Space will Terminate` has been displayed on the MVS system console.	An address space with the same name as the SNALINK LU6.2 address space is already active.	Check to see whether the address space with the same name is no longer required before stopping it, or rename the SNALINK LU6.2 address space. Restart the SNALINK LU6.2 address space.
The messages `Error 0000005A in VTAM OPEN` and `Errors detected in VTAM Initialization - Address Space will terminate` have been displayed on the MVS system console.	The SNALINK LU6.2 address space has not been able to find the VTAM application LU that has been defined in the VTAM statement of the SNALINK LU6.2 configuration data set.	This problem might be resolved by one or both of the following solutions: • Check the status of the SNALINK LU6.2 VTAM application LU and its VTAM major node. If it is not in a CONCT state, the VTAM major node and then the VTAM application LU must be activated. See "Useful VTAM operations" on page 581 for a detailed description of the VTAM operations that display the status of VTAM application LUs and activate them. • Check the VTAM application LU specified in the VTAM statement of the SNALINK LU6.2 configuration data set. Ensure that it exists and is not duplicated within the domain in which the SNALINK LU6.2 application program resides. See z/OS Communications Server: IP Configuration Reference for details about the SNALINK LU6.2 VTAM statement syntax and the SNALINK LU6.2 VTAM application LU definition.
The messages `Error 00000024 in VTAM OPEN` and `Errors detected in VTAM Initialization - Address Space will terminate` have been displayed on the MVS system console.	VTAM security is not allowing the SNALINK LU6.2 address space to access the VTAM application LU.	Check to see whether the SNALINK LU6.2 configuration data set VTAM statement password matches the password set in the VTAM application LU definition and correct it, if necessary. See z/OS Communications Server: IP Configuration Reference for details about the SNALINK LU6.2 VTAM statement syntax and the SNALINK LU6.2 VTAM application LU definition.
The SNALINK LU6.2 address space abends with a system abend code of 300 after the `SNALINK LU6.2 address space STARTED` message.	The abend code of 300 indicates that there is insufficient storage for the SNALINK LU6.2 address space.	Either increase the value of the REGION parameter for the address space or reduce the number of buffers specified in the SNALINK LU6.2 configuration data set. See z/OS Communications Server: SNA Messages and z/OS Communications Server: IP and SNA Codes for detailed SNALU6.2 abend code descriptions.

Table 43. Common SNALINK LU6.2 address space startup problems (continued)

If this is displayed. . .	Then this might have occurred . . .	Resolution
The SNALINK LU6.2 address space abends with an abend code of S0F8 after the `Initialization Complete...` message.	The MVS S0F8 abend code indicates that an SVC was issued in SRB mode. SNALINK LU6.2 is not designed to run with VTAM in SRB mode.	The SRBEXIT option in the VTAM application LU definition has been set to "Yes." Correct the VTAM application LU definition. See z/OS Communications Server: IP Configuration Reference for details about the SNALINK LU6.2 VTAM application LU definition.

If, after investigation, you do not find the SNALINK LU6.2 startup problem, obtain a description of all abend codes and errors written to the SYSPRINT data set and MVS system console. Most solutions to SNALINK LU6.2 address space starting problems can be solved by reading the error message or abend code descriptions.

See "Finding abend and sense code documentation" on page 589 and "Finding error message documentation" on page 589 for a list of references that contain SNALINK LU6.2 error message and abend code documentation.

DLC connection problems

These problems are related to the TCP/IP DLC connection between the TCP/IP address space and the SNALINK LU6.2 address space.

The DLC connection between the TCP/IP and SNALINK LU6.2 address spaces is established during the SNALINK LU6.2 address space startup after the SNALINK LU6.2 configuration data set has been parsed. This DLC connection can be established independently of the SNA LU type 6.2 connection between two or more SNALINK LU6.2 address spaces. The fundamental requirements of the DLC connection are an active, configured SNALINK LU6.2 address space and an active, configured TCP/IP address space. The DLC connection is initiated by a START statement in *hlq*.PROFILE.TCPIP.

Steps for checking DLC connection status

Check the status of the DLC connection.

Procedure

Perform the following steps:

1. Note the SNALINK LU6.2 address space startup messages displayed on the MVS system console.
2. Issue a NETSTAT DEVLINKS command to obtain the status of the DLC connection.

   See "Using NETSTAT" on page 580 for details about how to use the NETSTAT command to identify the status of the DLC connection.

Results

If the DLC connection status is not Connected, check the list of common DLC connection problems in the next section.

Analysis

Table 44 lists some of the common DLC connection problems between the SNALINK LU6.2 address space and the TCP/IP address space.

Table 44. Common DLC connection problems

If this is displayed. . .	Then this might have occurred . . .	Resolution
The message `Error in DLC connect...` has been displayed on the MVS system console and the NETSTAT DEVLINKS output shows that the DLC connection status is either `Issued Connect` or `Will retry connect`.	The TCP/IP address space is attempting to attach to the SNALINK LU6.2 address space, but the SNALINK LU6.2 address space is not responding.	Check whether the SNALINK LU6.2 address space is active and start it, if necessary.

Table 44. Common DLC connection problems (continued)

If this is displayed. . .	Then this might have occurred . . .	Resolution
The SNALINK LU6.2 address space has started, but the `Link open` message has not been displayed on the MVS system console, no other error messages have been displayed on the console, and the NETSTAT DEVLINKS output shows that the DLC connection status is either `Issued Connect` or `Will retry connect`.	This problem can be due to one of the following situations: 1. The SNALINK LU6.2 address space might be rejecting the connect attempt from the TCP/IP address space because it has the wrong TCP/IP ID. 2. The SNALINK LU6.2 address space might be rejecting the connect attempt from the TCP/IP address space because of a SNALINK LU6.2 link name that is incorrectly defined.	• Check the SNALINK LU6.2 SYSPRINT output for the "Rejecting DLC path for the link_name, wrong TCP/IP id tcpip_addr_space" error message.If this error message is displayed, check whether a valid TCPIP.DATA data set was specified as the SYSTCPD ddname in the SNALINK LU6.2 cataloged procedure and correct it, if necessary. **Note:** SYSTCPD can be overridden by the global TCPIP.DATA file. See z/OS Communications Server: IP Configuration Reference for an example of a SNALINK LU6.2 cataloged procedure and for the search order for the TCPIP.DATA data set. If a valid TCPIP.DATA data set has been used, check the TCP/IP address space specified in the TCPIPJOBNAME statement within it. See z/OS Communications Server: IP Configuration Reference for a detailed description of the TCPIPJOBNAME statement in the TCPIP.DATA. • Check the SNALINK LU6.2 SYSPRINT output for the "Rejecting DLC path for link_name, not configured" error message. If this error message is displayed, check to see whether the link name specified in the LINK statement of the SNALINK LU6.2 configuration data set matches the link name specified in the LINK statement associated with the SNALINK LU6.2 device defined in *hlq*.PROFILE.TCPIP. See z/OS Communications Server: IP Configuration Reference for details about the SNALINK LU6.2 LINK statement syntax and the TCPIP LINK statement syntax.

Table 44. Common DLC connection problems (continued)

If this is displayed. . .	Then this might have occurred . . .	Resolution
The SNALINK LU6.2 address space has been started but the `Link opened` message has not been displayed and the NETSTAT DEVLINKS output shows that the DLC connection is `Inactive`.	The DLC connection to the SNALINK LU6.2 device associated with the SNALINK LU6.2 address space might not have been started by the TCP/IP address space.	Check the START statements in *hlq*.PROFILE.TCPIP. If the SNALINK LU6.2 device has not been started, use the VARY TCPIP,*procname*,START,*device_name* for the SNALINK LU6.2 device or include the START statement in the *hlq*.PROFILE.TCPIP and restart the TCP/IP address space. See z/OS Communications Server: IP Configuration Reference for a detailed description of the START statement in the *hlq*.PROFILE.TCPIP.

Network connection establishment problems

These problems are related to the establishment of the SNA LU type 6.2 connection between two or more SNALINK LU6.2 devices.

The SNA LU type 6.2 connection can be established independently of the TCP/IP address space and the DLC link. The fundamental requirements for establishing the LU type 6.2 connection are two active, configured SNALINK LU6.2 devices that have an active SNA network connection between them.

Initiate the establishment of a network connection in one of the following ways:

- Connections with the INIT parameter specified on the DEST statement in the SNALINK LU6.2 configuration data set are established when the SNALINK LU6.2 address space is started.
- Connections with the DATA parameter specified on the DEST statement in the SNALINK LU6.2 configuration data set or connections that have timed out or been terminated are established when a request is made to the SNALINK LU6.2 address space to transfer data across the link.
- Connections can be established using the SNALINK LU6.2 RESTART MODIFY subcommand.

Steps for checking network connection problems

Issue MODIFY subcommands to the MVS SNALINK LU6.2 address space to check the status of the LU type 6.2 connection

Procedure

Perform the following steps:

1. MODIFY addr_sp_name,LIST,LU=dest_lu_name

   where addr_sp_name is the MVS SNALINK LU6.2 address space name and dest_lu_name is the SNA destination LU name of the remote SNALINK LU6.2 device.

   See “Using the SNALINK LU6.2 subcommand” on page 580 for more information about issuing this command and reading the output.

   If the connection status is not “Allocated,” continue with the following commands.

2. MODIFY addr_sp_name,RESTART,LU=dest_lu_name

   This command attempts to establish the LU type 6.2 connection between the SNALINK LU6.2 devices. During connection establishment, any problems causes error messages to be output to the MVS system console.
3. MODIFY addr_sp_name,LIST,LU=dest_lu_name

Results

If the connection status is still not "Allocated," note the messages in the SYSPRINT data set and on the MVS system console and continue with the following analysis.

Analysis

Table 45 lists some of the common SNALINK LU6.2 address space network establishment problems.

Table 45. SNALINK LU6.2 address space network establishment problems

If this is displayed. . .	Then this might have occurred . . .	Resolution
The SNALINK LU6.2 address space issued error message: `Unable to allocate send conversation`.	This problem can be due to one of the following situations: 1. The local VTAM application LU might not be enabled for LU type 6.2 conversations. The name of this LU is specified on the VTAM statement in the SNALINK LU6.2 configuration data set. 2. The remote VTAM application LU names might not identify an LU that is reachable or that can establish an LU type 6.2 conversation over the SNA network. The remote VTAM application LU name is specified in the DEST statement of the SNALINK LU6.2 configuration data set. For dependent LUs, both the SEND and RECV LU names must be able to establish LU type 6.2 conversations.	1. The APPC option in the VTAM application LU definition must be set to YES to enable LU type 6.2 conversations. 2. The first step is to check to see whether the remote SNALINK LU6.2 device is active. If the remote SNALINK LU6.2 is using VTAM to access the SNA network, see "Useful VTAM operations" on page 581 to check the active status of the remote LU. If the remote SNALINK LU6.2 device is active, use the VTAM error messages to determine why the LU type 6.2 conversation cannot be established with the destination LU. The VTAM error messages are written to the MVS system console immediately before the `Unable to allocate send conversation` message. VTAM sense code documentation can be found in z/OS Communications Server: SNA Messages and z/OS Communications Server: IP and SNA Codes. These messages might indicate a network configuration or hardware error.
VTAM error message output to the MVS system console: `REQUIRED LOGMODE NAME UNDEFINED`.	To allocate LU type 6.2 conversations over an SNA network, both sides must specify matching log modes. The VTAM log modes are defined in log mode tables. The log mode configured for use with this connection cannot be found in the log mode table specified on the VTAM application LU definition.	The log mode entry name specified as the LOGMODE parameter on the LINK statement in the SNALINK LU6.2 configuration file must exist in the log mode table specified on the MODETAB statement in the VTAM application LU definition.

The following list contains some of the common SNALINK LU6.2 address space network establishment problems. Each error symptom is listed with possible causes and resolutions.

Network connection loss problems

SNA network connection loss can be either expected or unexpected. This section deals with unexpected connection problems. The definitions of expected and unexpected losses are described before continuing with the analysis for unexpected loss.

Connections for the SNALINK LU6.2 address space can be configured to be normally active or normally inactive. The normally inactive configuration is used when there is a cost involved with the network connection time. Normally inactive connections are expected to experience connection establishment and loss regularly with use. Because of this, the SNALINK LU6.2 address space does not write messages to the MVS system console for connection loss. Connection loss for a normally active connection is unexpected. In this case, the SNALINK LU6.2 address space writes connection loss messages to the MVS system log.

When a connection is configured with the INIT parameter on the DEST statement and a timeout value of zero on the LINK statement in the SNALINK LU6.2 configuration data set, the connection is a normally active connection.

When a connection is configured with the DATA parameter on the DEST statement and a nonzero timeout value on the LINK statement in the SNALINK LU6.2 configuration data set, the connection is a normally inactive connection.

Check the connection that is experiencing the loss to ensure that the loss is unexpected. If the connection loss experienced is caused by errors, the loss is unexpected regardless of the connection configuration.

Documentation

Unexpected connection loss occurs if the SNALINK LU6.2 address space encounters errors that compromise the connection. In this case, error messages are written to the data set specified on the SYSPRINT DD statement in the SNALINK LU6.2 cataloged procedure.

To check the status of the SNA LU type 6.2 connection, issue the LIST MODIFY subcommand to the MVS SNALINK LU6.2 address space. See “Using the SNALINK LU6.2 subcommand” on page 580 for more information about issuing this command and reading the output.

Analysis

Use the error messages in the SNALINK LU6.2 SYSPRINT data set to identify the cause of the loss. See “Finding error message documentation” on page 589 for details on finding the documentation for these messages. Text in the message documentation specifies the action required to fix the problem.

Table 46 on page 577 lists an example of an outage problem.

Table 46. Outage problem

If this is displayed. . .	Then this might have occurred . . .	Resolution
The message EZA5797E `Rejecting bind from xxxxx-no DLC found`, along with VTAM error message IST663I `Bind fail request received, SENSE=800A0000`, is displayed on the MVS system console.	Large packet size sent in a PIU is rejected by the NCP with sense 800A0000 (PIU too large).	The PIU includes the TH, RH, and RU. SNALINK attempts to send data up to the MAXRU size. The total size of the PIU includes the RU portion and the additional 29 bytes for the TH and RH. If this exceeds the maximum size, NCP issues a negative response with sense 800A0000 (PIU too large), which results in the SNA session being taken down between SNALINK and the NCSTLU. When the DLC connection is reestablished, the NCP sends a Bind RU which is then rejected with sense 800A0000. The definitions used in the NCP and SNALINK must be such that MAXRU is at least 29 bytes less than MAXDATA. See z/OS Communications Server: SNA Network Implementation Guide for more information about defining the MAXDATA, MAXBFRU, and UNITSZ operands.

Data loss problems

These problems are related to data transfer over the SNALINK LU6.2 network. The first step is to determine the point in the network where the data is being lost. The following information is mainly concerned with determining the actual place of loss.

Steps for documenting data loss problems

Determine the source of the data loss.

Before you begin

To determine where the data packets are being lost, use the LIST MODIFY command for the SNALINK LU6.2 address space. See "Using the SNALINK LU6.2 subcommand" on page 580 for details. When listing the connection status, the number of packets sent and received over the connection since establishment is displayed in the report.

Procedure

Perform the following steps:

1. Record the current packet count for the SNALINK LU6.2 devices in the network that support the LIST MODIFY command.
2. Issue the PING command on one end of the connection. In a correctly functioning network, PING sends a data packet to the other end of the connection, which then sends a response data packet back to the PING command.
3. Use the updated packet counts to determine how far the packet went.
4. Issue the PING command from the other end of the connection.

5. Use the updated packet counts to determine how far the packet got.

Results

Tip: IP packet trace, as described in "Using IP packet trace" on page 587, can also be used to trace and validate the IP data packets as they enter and leave the SNALINK LU6.2 address space.

Analysis

Table 47 lists some of the common SNALINK LU6.2 data loss problems.

Table 47. Common SNALINK LU6.2 data loss problems

If this is displayed. . .	Then this might have occurred . . .	Resolution
Data packets are lost between the TCP/IP and the SNALINK LU6.2 address space (either end).	This problem can be due to one of the following situations: 1. The DLC link between the TCP/IP address space and the SNALINK LU6.2 address space might not be active. 2. The SNALINK LU6.2 address space might be discarding packets.	1. See "DLC connection problems" on page 571 to diagnose the DLC link problem. 2. When a condition occurs in the SNALINK LU6.2 address space that causes data to be lost, "discarding datagram" messages are written to the data set specified by the SYSPRINT DD statement in the SNALINK LU6.2 cataloged procedure. See "Finding error message documentation" on page 589 for details on finding the documentation for these messages. Text in the message documentation specifies the action required to fix the problem.
Data packets are actually not lost but the protocol (PING) times out.	The SNALINK LU6.2 device might be establishing the LU type 6.2 connection to transfer the data packets. The delay in establishing the connection might be causing the protocol to time out.	If the DATA parameter is specified on the DEST statement for the connection in the SNALINK LU6.2 configuration data set, the connection is not established until data is to be transferred over the connection. In this case, after the first data transfer, further data packets are transferred successfully. If the TIMEOUT parameter is specified on the LINK statement for the connection in the SNALINK LU6.2 configuration data set, the connection can be timing out too often, causing the connection to be reestablished for each data transfer. In this case, the protocol timeout value or the connection timeout value should be increased.
Data packets are lost between the SNALINK LU6.2 devices.	The network is failing.	Check for VTAM error messages on the MVS system console. See "VTAM buffer traces" on page 589 for more details about using VTAM traces to diagnose the SNA network.

Data corruption problems

To determine the source of corruption for the data packets, use the IP packet tracing facility. This facility traces and validates the IP data packets as they enter and leave the SNALINK LU6.2 address space. Using this facility, the source of corruption can be identified as either the SNA network or the TCP/IP system.

Documentation

Set up the network conditions that are experiencing the data corruption. Start component trace in the SNALINK LU6.2 address space. Use the appropriate amount of data and time to ensure the corruption occurs.

Guideline: Allocate the MVS GTF trace data set (usually SYS1.TRACE) large enough to hold the expected trace output. This trace data set wraps back to the start of the data set when full, overwriting trace information. When tracing, this option does not collect all the data, which means the corruption could be missed. When formatting, this option turns off some of the IP packet validation processing.

Analysis

The IP packet trace facility analyzes the data corruption problem automatically. After the trace is collected, the trace data is passed through a formatter, which presents the data packets in an easy-to-read report and validates the contents of the packets against the RFC requirements. Every byte of the data packet is validated including reserved fields. The checksums are also recalculated and verified. If any of the data packets traced are corrupted, the formatter writes messages in the formatted report.

You can use this method, possibly together with TCP/IP internal traces, network level traces, or both, to identify the source and type of corruption.

For details on how to use the IP packet trace facility, see “Using IP packet trace” on page 587.

Documentation references for problem diagnosis

This section contains the information and documentation references required to gather and decode diagnostic information about the SNALINK LU6.2 network connection.

The main tools used for problem diagnosis are the NETSTAT utility, the SNALINK LU6.2 LIST subcommand, VTAM status display operations, the SNALINK LU6.2 internal trace facility, and the IP packet trace facility. The use of these tools is explained in the following sections. An explanation of how to interpret the output from each of these tools is also provided and referenced against the sample output.

For TCP/IP internal tracing or VTAM buffer tracing, you are referred to the appropriate diagnosis documentation.

Two cross-reference sections are provided at the end of this section that list all of the types of abend codes, sense codes, and error messages that can be issued from the SNALINK LU6.2 network connection. For each type of abend code, sense code, or error message, you are referred to the documentation that provides a complete description.

Using NETSTAT

This section describes how to use NETSTAT to query the state of TCP/IP devices. This command can be used to quickly verify the status of the SNALINK LU6.2 device and link with relation to the TCP/IP address space.

The NETSTAT DEVLINKS command output displays only information that is known to TCP/IP.

Restriction: The TCP/IP address space must be started before the NETSTAT command can query the connection status.

The NETSTAT DEVLINKS command displays the devices and links that have been defined to the main TCP/IP address space and the status of these devices (whether active or inactive).

Figure 67 shows a sample of output from the NETSTAT DEVLINKS command.

```
DevName: SNALU621          DevType: SNALU62
  DevStatus: Ready
  LnkName: SNALU62L          LnkType: SNALU62     LnkStatus: Ready
    NetNum: 1   QueSize: 0
    BytesIn: 0                     BytesOut: 0
    ActMTU: 32764
  BSD Routing Parameters:
    MTU Size: 00000             Metric: 00
    DestAddr: 0.0.0.0           SubnetMask: 255.0.0.0
  Multicast Specific:
    Multicast Capability: No
```

Figure 67. NETSTAT DEVLINKS output example

The example shows four SNALINK LU6.2 devices and associated links known to TCP/IP.

The most significant field for diagnosing DLC connection problems is the DevStatus field. See z/OS Communications Server: IP Configuration Reference for detailed interpretation of the device status and its importance in the SNALINK LU6.2 DLC connection.

Using the SNALINK LU6.2 subcommand

This section details how to use the LIST MODIFY subcommand for the MVS SNALINK LU6.2 address space. The SNALINK LU6.2 address space has interactive commands to control the operation and list the status of the active address space. The LIST MODIFY subcommand writes a report to the MVS system console giving the status of the specified connections.

The connection status listed by the LIST subcommand can be requested for a particular remote VTAM application LU name or destination IP address. The following is an example using the LU parameter:

```
MODIFY procname,LIST LU=lu_name
```

In this example, *procname* is the member name of the cataloged procedure used to start the local SNALINK LU6.2 address space and lu_name is the remote VTAM application LU name of the connection for which you are requesting the status.

Figure 68 on page 581 shows a sample output from the subcommand.

```
f snal621a,list lu=snal622a
EZA5971I LIST ACCEPTED; RANGE = SINGLE CONNECTION
EZA5967I   9.67.22.2 (Connected on 01.051 at 15:44:32)
EZA5968I     Connected via:  DATA            Trace Level: ON
EZA5969I     SEND:-  Status: Allocated       Packets Out: 1
EZA5970I     RECV:-  Status: Allocated       Packets In:  1
EZA5974I LIST COMPLETED
```

Figure 68. LIST MODIFY subcommand output example

An active connection displays the EZA5968I `Connected` message with the `Allocated` status for both the send and receive conversations.

The SNALINK LU6.2 connection allocates two LU type 6.2 conversations: one for sending data to the remote device and one for receiving data. For independent LUs, the remote LU name is the same for both conversations. For dependent LUs, a remote LU name is specified for both the send and receive conversations.

The `Packets In` and `Packets Out` fields are decimal counters that record the number of data packets received from the remote SNALINK LU6.2 and the number of data packets sent to the remote SNALINK LU6.2, respectively. These fields can be used to identify configuration errors that cause data packets to be lost or discarded. For example, the packet counters can be used to track how far a PING packet travels around the network circuit before it gets lost. Each counter incremented means the packet made it past that point.

For more information about the contents of the messages from the LIST MODIFY subcommand, see the message documentation referenced in "Finding error message documentation" on page 589. See the z/OS Communications Server: IP System Administrator's Commands for more explanation of the LIST MODIFY subcommand.

Useful VTAM operations

This section describes how to use the VTAM DISPLAY and VARY commands to activate an LU, change an LU definition, and to check the status of an application LU.

VTAM application LUs are defined with VTAM macros in a member of the SYS1.VTAMLST data set. The data set member, called the major node, can contain many application LU definitions, called minor nodes. The application LU names (minor node names) are specified on the VTAM and DEST statements in the SNALINK LU6.2 configuration data set.

Activating an LU

To activate an LU, the major node containing the LU definition must be activated first. If there are no definition errors, all the minor nodes defined in the major node are activated when the major node is activated. If a minor node becomes inactive, it can be activated individually. The following is an example of a VTAM VARY subcommand to activate a major or minor node:

```
VARY NET,ACT,ID=node_name
```

In this example, *node_name* is the major or minor node name to activate.

See "Displaying the status of an LU" on page 582 for an explanation of the active states for a minor node.

See z/OS Communications Server: SNA Operation for a complete description of the VARY ACT subcommand.

Changing an LU definition

To change an LU (minor node) definition, the major node containing the LU definition must be deactivated and then reactivated to force VTAM to read the new definition. The following is an example of a VTAM VARY subcommand to deactivate a major node:

```
VARY NET,INACT,ID=majnode_name
```

In this example, *majnode_name* is the major node name to deactivate.

See "Activating an LU" on page 581 for the major node activation subcommand.

See z/OS Communications Server: SNA Operation for a complete description of the VARY INACT subcommand.

Displaying the status of an LU

To display the status of an LU definition, use the following command:

```
DISPLAY NET,ID=node_name,E
```

In this example, *node_name* is the major or minor node name for which you want to display the status.

Displaying the status of a major node lists all of the minor nodes defined to the major node and their STATUS field. For complete information about status of a minor node, specify the actual minor node name in the command.

The STATUS field for a successfully activated LU definition is set to "CONCT," which means connectable. An LU in this state is waiting for the SNALINK LU6.2 address space to be started. An LU in the CONCT state cannot establish an LU type 6.2 conversation.

Figure 69 on page 583 shows a sample of the output from an LU in connectable state.

```
D NET,ID=SNAL621A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = NETA.SNAL621A, TYPE = APPL 573
IST486I STATUS= CONCT, DESIRED STATE= CONCT
IST1447I REGISTRATION TYPE = CDSERVR
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=LU62MODE USS LANGTAB=***NA***
IST1632I VPACING =  0
IST597I CAPABILITY-PLU INHIBITED,SLU INHIBITED,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = SNALNK1A
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = ***NA***, STEPNAME = ***NA***, DSPNAME = ***NA***
IST228I ENCRYPTION = OPTIONAL , TYPE = DES
IST1563I CKEYNAME = SNAL621A CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = ***NA*** MAXIMUM = ***NA***
IST171I ACTIVE SESSIONS = 0000000000, SESSION REQUESTS = 0000000000
IST172I NO SESSIONS EXIST
IST314I END
```

Figure 69. DISPLAY subcommand output example for connectable LU

After the SNALINK LU6.2 address space has successfully started, the STATUS field is set to ACTIV, which means in use by an address space.

Figure 70 shows a sample of the output from a DISPLAY command for an LU in active state.

```
D NET,ID=SNAL621A,E
IST097I DISPLAY ACCEPTED
IST075I NAME = NETA.SNAL621A, TYPE = APPL 545
IST486I STATUS= ACT/S, DESIRED STATE= ACTIV
IST1447I REGISTRATION TYPE = CDSERVR
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=LU62MODE USS LANGTAB=***NA***
IST1632I VPACING =  0
IST597I CAPABILITY-PLU ENABLED  ,SLU ENABLED  ,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = SNALNK1A
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = SNAL621A, STEPNAME = SNAL621A, DSPNAME = IST84E76
IST228I ENCRYPTION = OPTIONAL , TYPE = DES
IST1563I CKEYNAME = SNAL621A CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = 0 MAXIMUM = 136
IST171I ACTIVE SESSIONS = 0000000003, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME     STATUS         SID          SEND RECV VR TP NETID
IST635I SNAL622A ACTIV-S    EAABEEC3D9B90C37 0002 0000  0  0 NETA
IST635I SNAL622A ACTIV/SV-S EAABEEC3D9B90C35 0002 0002  0  0 NETA
IST635I SNAL622A ACTIV-P    F6ABEEC3DCB90B7D 0000 0002  0  0 NETA
IST314I END
```

Figure 70. DISPLAY subcommand output example for active LU

This example shows that the SNALINK LU6.2 address space (SNAL621A) has been started successfully and has its local LU (SNAL621A) in use with three sessions active to a remote LU (SNAL622A).

For each SNALINK LU6.2 connection, VTAM establishes three sessions between the application LUs. The first is the control session, which is the middle session in the example. The other two sessions are established for the LU type 6.2 conversations allocated for the connection, one for sending data and one for receiving data.

See z/OS Communications Server: SNA Operation for more information about the DISPLAY command.

Traces

Use the following traces to obtain information about the data flows and actions of the SNALINK LU6.2 network connection:

- SNALINK LU6.2 internal trace
- IP packet trace
- TCP/IP internal trace
- VTAM buffer trace

The SNALINK LU6.2 internal trace is the most useful for determining the state of the SNALINK LU6.2 address space. The IP packet trace facility is the most helpful trace facility for monitoring IP packets transferred across the SNALINK LU6.2 network. The TCP/IP internal traces can be used to diagnose problems with the DLC link between TCP/IP and SNALINK LU6.2. The VTAM buffer trace is used to monitor data transactions through the VTAM API interface.

Using SNALINK LU6.2 internal traces

The SNALINK LU6.2 internal traces are written to the location specified by the SYSPRINT statement in the SNALINK LU6.2 cataloged procedure. These traces provide information about the internals of the SNALINK LU6.2 address space.

SNALINK LU6.2 internal tracing is enabled by specifying the following statement in the SNALINK LU6.2 configuration data set:

```
TRACE DETAIL ALL
```

The SNALINK LU6.2 internal trace can also be started by passing a MODIFY console command to the SNALINK LU6.2 interface. The following MODIFY command starts internal tracing:

```
MODIFY procname,TRACE DETAIL ALL
```

In this example, *procname* is the member name of the cataloged procedure used to start the local SNALINK LU6.2 address space.

See the z/OS Communications Server: IP Configuration Reference for detailed descriptions of the TRACE statement parameters and the TRACE subcommand parameters.

```
EZA5926I LU62CFG : STARTING PASS 1 OF 2
EZA5926I LU62CFG : STARTING PASS 2 OF 2
 1 EZA5927I LU62CFG : NO ERRORS DETECTED - INITIALIZATION WILL CONTINUE
EZA5932I INITIALIZATION COMPLETE - APPLID: SNAL621A  TCP/IP: TCPCS
EZA5997I CONNECTION 9.67.22.2 WILL TIMEOUT IN 600 SECONDS
 7 EZA5994I CNOS FOR INDEPENDENT PARTNER; SESSLIM=0002, WINNER=0001, LOSER=0001
EZA5984I   OLU= SNAL621A, DLU= SNAL622A, IP ADDRESS= 9.67.22.2
EZA6029E OPRCNOS  ERR. R15 00000000 R0 0000000B RTNCD 00000000 FDBK2 0000000B
EZA6030E OPRCNOS  ERR.  RCPRI= 0008,  RCSEC= 0001
EZA6031E OPRCNOS  SENSE CODE RECEIVED:  08570003
EZA6032E          SENSE CODE SPECIFIED: 00000000
EZA6023E UNABLE TO COMPLETE CNOS ON LU SNAL622A FOR 9.67.22.2
 2 EZA6009W CONVERSATIONS FOR 9.67.22.2 TERMINATED
EZA6011E UNABLE TO ALLOCATE SEND CONVERSATION FOR 9.67.22.2
 1 EZA5933I LINK SNALU62L OPENED
EZA5997I CONNECTION 9.67.22.2 WILL TIMEOUT IN 600 SECONDS
EZA5936I RECEIVE CONVERSATION ALLOCATED FOR 9.67.22.2
EZA5986I VTAM CONVERSATION ALLOCATED; CONVID= 01000011, SID= F6ABEEC3DCB90B89
EZA5984I   OLU= SNAL622A, DLU= SNAL621A, IP ADDRESS= 9.67.22.2
EZA5997I CONNECTION 9.67.22.2 WILL TIMEOUT IN 600 SECONDS
 2 EZA5935I SEND CONVERSATION ALLOCATED FOR 9.67.22.2
EZA5986I VTAM CONVERSATION ALLOCATED; CONVID= 01000012, SID= EAABEEC3D9B90C3F
EZA5984I   OLU= SNAL621A, DLU= SNAL622A, IP ADDRESS= 9.67.22.2
EZA5997I CONNECTION 9.67.22.2 WILL TIMEOUT IN 600 SECONDS
 4 EZA5992I IP DATAGRAM ADDED TO THE VTAM SEND QUEUE, LENGTH= 276,
QUEUE COUNT = 1
EZA5985I   LU= SNAL622A, LINKNAME= SNALU62L, IP ADDRESS = 9.67.22.2
 3 EZA5988I VTAM SENT LOGICAL RECORD; CONVID= 01000012,
SID= EAABEEC3D9B90C3F, LENGTH= 280
EZA5984I   OLU= SNAL621A, DLU= SNAL622A, IP ADDRESS= 9.67.22.2
EZA5995I NUMBER OF IP PACKETS SENT ON 9.67.22.2 = 1
 6 EZA5999I   450001140315000040014D4C094316010943010208 00D58BEFBF74
0AB56D2E96EF8F85CA03EFDD80
EZA5999I   78B6FE8035FE858058EE968017DAAD8015B4AE80092C55803DA14680271B7D807DCC5
E8074502580
EZA5999I   39CDF68000F24D8023BE0E8012A9F5805434A6804C9F1D806249BE805779C5807B955
6800961ED80
EZA5999I   7C2F6E800DFF958006B006800E7ABD801C2F1E80597B65803444B6800B298D805508
CE80352D3580
EZA5999I   2B13668006AE5D80217C7E807455058079DC168060492D80644A2E804232D580175E
C6804F39FD80
EZA5999I   6831DE802206A580625B768062C0CD805FF38E806F10758021922680021D9D80664F3
E805C904580
EZA5999I   03C2D6806C906D807E04EE8075C615801FAD868039593D8011D49E801DF1E5807412
368000000000
EZA5997I CONNECTION 9.67.22.2 WILL TIMEOUT IN 600 SECONDS
EZA5997I CONNECTION 9.67.22.2 WILL TIMEOUT IN 600 SECONDS
 3 EZA5989I VTAM RECEIVED LOGICAL RECORD; CONVID= 01000011,
SID= F6ABEEC3DCB90B89, LENGTH= 280
EZA5984I   OLU= SNAL622A, DLU= SNAL621A, IP ADDRESS= 9.67.22.2
EZA5996I NUMBER OF IP PACKETS RECEIVED ON 9.67.22.2 = 1
 6 EZA5999I   4500011402E8000040014D790943010209431601 0000DD8BEFBF74
0AB56D2E96EF8F85CA03EFDD80
EZA5999I   78B6FE8035FE858058EE968017DAAD8015B4AE80092C55803DA14680271B7D807DCC5
E8074502580
EZA5999I   39CDF68000F24D8023BE0E8012A9F5805434A6804C9F1D806249BE805779C5807B955
6800961ED80
EZA5999I   7C2F6E800DFF958006B006800E7ABD801C2F1E80597B65803444B6800B298D805508C
E80352D3580
EZA5999I   2B13668006AE5D80217C7E807455058079DC168060492D80644A2E804232D580175E
C6804F39FD80
EZA5999I   6831DE802206A580625B768062C0CD805FF38E806F10758021922680021D9D80664F3
E805C904580
```

Figure 71. SNALINK LU6.2 internal trace output

```
EZA5999I   03C2D6806C906D807E04EE8075C615801FAD868039593D8011D49E801DF1E5807412
368000000000
 5 EZA5990I IP DATAGRAM PACKED INTO MESSAGE, LENGTH= 276
EZA5985I   LU= SNAL622A, LINKNAME= SNALU62L, IP ADDRESS = 9.67.22.2
EZA5938I RECEIVED OPERATOR SHUTDOWN REQUEST
 2 EZA5937I LINK SNALU62L CLOSED
 3 EZA5989I VTAM RECEIVED LOGICAL RECORD; CONVID= 01000011,
SID= F6ABEEC3DCB90B89, LENGTH= 280
EZA5984I   OLU= SNAL622A, DLU= SNAL621A, IP ADDRESS= 9.67.22.2
EZA5996I NUMBER OF IP PACKETS RECEIVED ON 9.67.22.2 = 1
 6 EZA5999I   4500011402E8000040014D790943010209431601000DD8BEFBF74
0AB56D2E96EF8F85CA03EFDD80
EZA5999I   78B6FE8035FE858058EE968017DAAD8015B4AE80092C55803DA14680271B7D807DCC5
E8074502580
EZA5999I   39CDF68000F24D8023BE0E8012A9F5805434A6804C9F1D806249BE805779C5807B955
6800961ED80
EZA5999I   7C2F6E800DFF958006B006800E7ABD801C2F1E80597B65803444B6800B298D805508C
E80352D3580
EZA5999I   2B13668006AE5D80217C7E807455058079DC168060492D80644A2E804232D580175E
C6804F39FD80
EZA5999I   6831DE802206A580625B768062C0CD805FF38E806F10758021922680021D9D80664F3
E805C904580
EZA5999I   03C2D6806C906D807E04EE8075C615801FAD868039593D8011D49E801DF1E5807412
368000000000
 5 EZA5990I IP DATAGRAM PACKED INTO MESSAGE, LENGTH= 276
EZA5985I   LU= SNAL622A, LINKNAME= SNALU62L, IP ADDRESS = 9.67.22.2
EZA5938I RECEIVED OPERATOR SHUTDOWN REQUEST
 2 EZA5937I LINK SNALU62L CLOSED
```

Following are brief explanations of the numbered items in the output:

1 Messages written to the MVS system console.

2 VTAM send and receive conversation status.

3 Information about the VTAM API interface data flow.

The VTAM interface information contains the LU type 6.2 conversation ID (Convid), the VTAM session ID (SID), length of the VTAM logical record, the origin and destination VTAM application LUs, and the home IP address.

The VTAM logical record length should be four greater than the length of the TCP/IP datagram packet to account for the VTAM logical record header.

4 Information about data received from the TCP/IP DLC connection.

Datagrams received from the SNALINK LU6.2 DLC connection are unpacked from the DLC message and added to the appropriate VTAM send queue for transmission.

5 Information about data received from the VTAM API interface.

Datagrams received from the VTAM API are packed into a DLC message buffer.

6 Hexadecimal display of data passed through the SNALINK LU6.2 address space.

There should be a hexadecimal display for every **4** and **5** event.

7 Change number of sessions (CNOS) data.

See the z/OS Communications Server: SNA Programmer's LU 6.2 Guide for more information about CNOS processing.

Using IP packet trace

Trace on the SNALU62 LINKNAME using the TCPIP PKTTRACE command, or on the SNA LU name using the VTAM **buffer trace** command.

If the LINKNAME parameter of the IP packet trace facility is specified, only packets transferred along the given link are traced. Specifying this parameter is recommended to avoid tracing many unrelated packets.

See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for details about how to use the IP packet trace facility.

Figure 72 on page 588 shows an example of a CTRACE formatted packet trace record.

```
   19 VIC127   PACKET   00000001 22:52:18.648744 Packet Trace
 To Link           : SNALU62L         Device: SNA_LU6.2        Full=276
  Tod Clock        : 2001/02/20 22:52:18.648743
  Lost Records     : 0                Flags: Pkt Ver2 Out
  Source Port      : 0                Dest Port: 0     Asid: 01F6 TCB: 007AEE88
 IpHeader: Version : 4                Header Length: 20
  Tos              : 00 QOS: Routine  Normal Service
  Packet Length    : 276              ID Number: 1543
  Fragment         :                  Offset: 0
  TTL              : 64               Protocol: ICMP           CheckSum: 4A5A FFFF
  Source           : 9.67.22.1
  Destination      : 9.67.1.2

 ICMP
  Type/Code        : ECHO             CheckSum: 5B3F FFFF
  Id               : 4923             Seq: 11849
  Echo Data        : 248
000000 B56D52DB 47A07ACA 2D523F5F 72BE75F9 |.mR.G.z.-R?_r.u.
000010 36332E6F 5A2D7969 5F7DDC7F 401715D9 |63.oZ-yi_}..@...
000020 2B9B598F 6414BB49 0D13B59F 08F8D9B9 |+.Y.d..I........
000030 099E00AF 643EE129 5C304ABF 667B4199 |....d>.)\0J.f{A.
000040 260FA3CF 4CCB6B09 3EE01BDF 6B45CD79 |&...L.k.>...kE.y
000050 33B4C2EF 2069D8E9 051FA8FF 618FFD59 |3... i......a..Y
000060 3441DE0F 5059AAC9 2EDB721F 49215139 |4A..PY....r.I!Q9
000070 2A5B752F 5A6A60A9 7DEFF73F 15514919 |*.u/Zj`.}..?.QI.
000080 0B96084F 26FB7A89 4829B85F 2B0764F9 |...O&.z.H)._+.d.
000090 7276176F 26FC7869 0945357F 1EBB24D9 |rv.o&.xi.E5...$.
0000A0 1070228F 31ECDA49 34EEEE9F 327408B9 |.p".1..I4...2t..
0000B0 5FE8A9AF 23DC2029 48C363BF 13C99099 |_...#. )H.c.....
0000C0 16342CCF 3B69CA09 1E4F14DF 59E33C79 |.4,.;i...O..Y.<y
0000D0 55972BEF 37C557E9 7D0E81FF 43788C59 |U.+.7.W.}...Cx.Y
0000E0 1F46270F 36AE49C9 6C6E2B1F 34D10039 |.F'.6.I.ln+.4..9
0000F0 05659E2F 52741FA9                   |.e./Rt..

IP Header          : 20     IP: 9.67.22.1, 9.67.1.2
000000 45000114 06070000 40014A5A 09431601  09430102

Data               : 256    Data Length: 256
000000 08005B3F 49232E49 B56D52DB 47A07ACA |..$....._....:. ...?I#.I.mR.G.z.|
000010 2D523F5F 72BE75F9 36332E6F 5A2D7969 |...¼...9...?!... -R?_r.u.63.oZ-yi|
000020 5F7DDC7F 401715D9 2B9B598F 6414BB49 |¼'." ..R........ _}..@...+.Y.d..I|
000030 0D13B59F 08F8D9B9 099E00AF 643EE129 |.....8R......... ............d>.)|
000040 5C304ABF 667B4199 260FA3CF 4CCB6B09 |*....#.r..t.<.,. \0J.f{A.&...L.k.|
000050 3EE01BDF 6B45CD79 33B4C2EF 2069D8E9 |....,.....B...QZ >...kE.y3... i..|
000060 051FA8FF 618FFD59 3441DE0F 5059AAC9 |..y./..........I ....a..Y4A..PY..|
000070 2EDB721F 49215139 2A5B752F 5A6A60A9 |.........$..!.-z ..r.I!Q9*.u/Zj`.|
000080 7DEFF73F 15514919 0B96084F 26FB7A89 |'.7......o.|..:i }..?.QI....O&.z.|
000090 4829B85F 2B0764F9 7276176F 26FC7869 |...¼...9...?.... H)._+.d.rv.o&.xi|
0000A0 0945357F 1EBB24D9 1070228F 31ECDA49 |..."...R........ .E5...$..p".1..I|
0000B0 34EEEE9F 327408B9 5FE8A9AF 23DC2029 |........¼Yz..... 4...2t.._...#. )|
0000C0 48C363BF 13C99099 16342CCF 3B69CA09 |.C...I.r........ H.c......4,.;i..|
0000D0 1E4F14DF 59E33C79 55972BEF 37C557E9 |.|...T...p...E.Z .O..Y.<yU.+.7.W.|
0000E0 7D0E81FF 43788C59 1F46270F 36AE49C9 |'.a.............I }...Cx.Y.F'.6.I.|
0000F0 6C6E2B1F 34D10039 05659E2F 52741FA9 |%>...J.........z ln+.4..9.e./Rt..|
```

Figure 72. A CTRACE formatted packet trace record

TCP/IP internal traces

The TCP/IP internal traces are written to the data set that is specified on the TCP/IP address space SYSDEBUG ddname statement. These traces provide information about the internals of the TCP/IP address space that can be used to diagnose problems in establishing the DLC link between the TCP/IP address space and the SNALINK LU6.2 address space.

VTAM buffer traces

The VTAM buffer traces provide information about the contents of the VTAM API buffers. This information can be used to follow the data through the VTAM API interface. For details about VTAM buffer tracing and reading the trace reports, see z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures.

Finding abend and sense code documentation

The following list refers to the appropriate abend and sense code documentation for all abend and sense codes expected in the SNALINK LU6.2 network connection:

- See z/OS Communications Server: IP Messages Volume 1 (EZA) and z/OS Communications Server: IP and SNA Codes for detailed SNALINK LU6.2 abend code descriptions.
- Sense codes in SNALINK LU6.2 error messages are generated by VTAM. See z/OS Communications Server: SNA Messages and z/OS Communications Server: IP and SNA Codes for detailed sense code descriptions.

Finding error message documentation

The following list refers to the appropriate error message documentation for all error messages expected when using SNALINK LU6.2:

- Error messages from SNALINK LU6.2 are written to the SNALINK LU6.2 SYSPRINT data set and the MVS system console. See z/OS Communications Server: IP Messages Volume 1 (EZA) and z/OS Communications Server: IP and SNA Codes for descriptions of the SNALINK LU6.2 error messages.
- Error messages from TCP/IP are written to the TCPIP SYSERROR data set. See z/OS Communications Server: IP Messages Volume 1 (EZA) and z/OS Communications Server: IP and SNA Codes for descriptions of the error messages in these data sets.
- Error messages from VTAM are written to the MVS system console. See z/OS Communications Server: SNA Messages and z/OS Communications Server: IP and SNA Codes for descriptions of the VTAM error messages written to the MVS system console.

Chapter 21. Diagnosing REXEC, REXECD, and RSH problems

This topic contains diagnosis information about the classic (non-z/OS UNIX) Remote Execution Protocol (REXEC), the Remote Execution Protocol Daemon (REXECD), and the remote shell client (RSH). See "General information about REXEC and RSH" for information about REXEC and RSH and "General information about REXECD" on page 592 for information about REXECD.

The following subtopics are included:

- "General information about REXEC and RSH"
- "General information about REXECD" on page 592

General information about REXEC and RSH

REXEC and RSH are remote execution clients that allow you to execute a command on a remote host and receive the results on the local host. REXEC and RSH commands can be executed from the TSO command line or as a batch program.

See z/OS Communications Server: IP Configuration Reference for information about defining the remote execution server.

Figure 73 shows the principle behind REXECD.

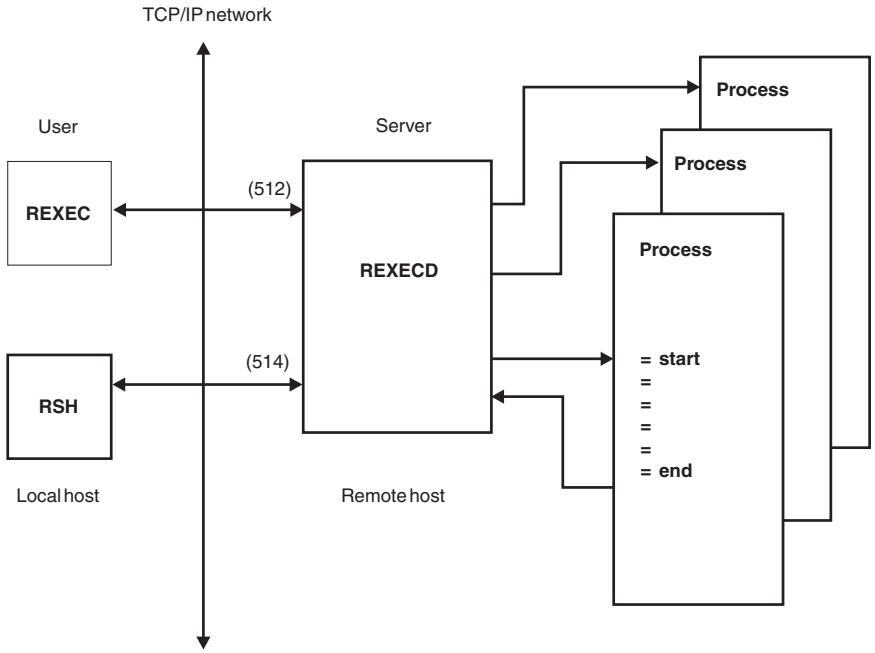

Figure 73. Remote execution protocol principle

Documentation for REXEC problem diagnosis

The following kinds of information might be required to diagnose a REXEC problem:

© Copyright IBM Corp. 2000, 2015

- REXEC console log
- REXEC debug trace

TSO console log

The TSO console log should be saved and made available, particularly if there are any error messages that are displayed at the console.

Activating the REXEC debug trace

To activate the REXEC debug trace, use the REXEC -d command.

See z/OS Communications Server: IP User's Guide and Commands for more information about REXEC commands.

REXEC trace example and explanation

Figure 74 shows an example of an REXEC trace. Short descriptions of the numbered items in the trace follow the figure.

REXEC trace output is sent to the TSO console from which the command was submitted.

```
rexec -d -l debfox -p mypwd norway time
Established affinity with TCPCS
EZA4801I MVS TCP/IP REXEC CS V1R5
EZA4775I Calling function rexec_af with the following:
Host: norway   user: debfox   cmd: time   port: 512
EZA4774I  rexec invoked;
Data socket = 1  Control socket = 3
IKJ56650I TIME-01:22:00 PM. CPU-00:00:00 SERVICE-5982 SESSION-00:00:01 March 24, 2003
EZA4789I  rexec complete
```

Figure 74. Example of an REXEC trace

RSH trace example and explanation

Figure 75 shows an example of an RSH trace. Short descriptions of numbered items in the trace follow the figure.

RSH trace output is sent to the RSH console.

```
rsh -d -l user1/tcpsup norway time
Established affinity with TCPCS
EZA5025I Calling function rcmd_af with the following:
Host: norway   user: user1   cmd: time   port: 514
EZA5046I  rsh invoked;
Data socket = 1  Control socket = 3
IKJ56650I TIME-02:30:30 PM. CPU-00:00:00 SERVICE-6454 SESSION-00:00:00 March 24,2003
EZA5048I  rsh complete
```

Figure 75. Example of an RSH trace

General information about REXECD

The remote execution server allows execution of a TSO batch command that is received from a remote host. REXECD supports both the remote execution command (REXEC) and remote shell (RSH) client protocols.

Note: When the REXECD server is active, it has outstanding listens on Ports 512 and 514. If you want to have a concurrent server for the z/OS UNIX REXECD or RSHD daemons, then configure them to use different ports.

Documentation for REXECD problem diagnosis

The following kinds of information might be required to diagnose a REXECD problem:

- REXECD console log
- REXECD traces

MVS system console log

Save the MVS system console log and make it available, particularly if there are any error messages displayed at the console.

Starting REXECD server traces

To run the REXECD trace, REXECD must be started with one or more of the following options on the TRACE parameter in the PROC statement:

LOG
Specifies to write trace records to the SYSPRINT data set.

SEND
Specifies to send trace records to the REXEC or RSH client.

CLIENT
Specifies a specific client host for which trace records are to be produced.

ALLCLIENTS
Specifies that host records are to be produced for all clients.

See the z/OS Communications Server: IP Configuration Reference for more information about the options. See the z/OS MVS JCL Reference for information about the length limit of the PARM= parameter on the exec statement in the start procedure.REXECD trace output is included in the job output log.

Restriction: If more than one trace option is selected, the options must be enclosed within parentheses.

Example of an REXECD trace of a client using the SEND command

Figure 76 on page 594 shows a portion of an example of an REXECD trace of a client using a SEND command. Short descriptions of numbered items in the trace follow the figure.

```
EZA4801I MVS TCP/IP REXEC CS V1R10
 1
EZA4437I SSCSARAY[0] JOB00043 40 WAITING
 2
EZA4436I SSCSARAY[0] JOB00043 80 ACTIVE
EZA4436I SSCSARAY[0] JOB00043 80 ACTIVE
 3
EZA4438I SSCSARAY[0] JOB00043 20 COMPLETED
EZA4385I SSSORT(CTRL): 00000000
 4
EZA4392I S99ret: 00000000, A RSHD NEWALTON.RSHD5.JOB00043.D0000105.?
 5
TIME-12:04:12 PM. CPU-00:00:00 SERVICE-1157 SESSION-00:00:01 APRIL 1,2008
EZA4393I S99ret: 00000000
 6
EZA4442I SSSORT(next): 00000004 NO MORE OUTPUT ON JES SPOOL
```

Figure 76. Example of an REXECD trace of a client using a SEND command

Following are short descriptions of the numbered items in the trace:

1 JOB00043 is the JES job number. The number 40 indicates that the job is waiting for execution. This means that the remote execution server has processed the REXEC client request, created a JES job, and has submitted the JOB to JES. The server continues to check the status.

Guidelines: If the status does not change from 40, this could indicate one of the following problems:

- A JES initiator has not started to process the submitted job class.
- Other jobs might be running in this class that are inhibiting this job from starting.

Make a note of the job number and check JES activity, or check any other jobs that might be running at the same time.

2 The number 80 indicates that the job is currently active. This means that the remote execution server has checked with JES on the job status and was informed that the job is running.

Guideline: If the status does not change from 80, this could indicate that the job is taking too long to run. REXEC was not intended to be used for long running jobs. Long running jobs should be submitted using FTP JES processing.

Tip: The FTP JES interface provides an alternate method of submitting a batch job remotely and retrieving the results.

3 The number 20 indicates that the job is on the output queue. This means that when the server checked with JES on the job status it discovered that the JES job has completed, and it has been placed on the output queue. The command output should be sent back to the client soon. See z/OS Communications Server: IP Messages Volume 1 (EZA) for more information about the individual messages in the trace.

4 This line shows the return code from the dynamic allocation of the JES data sent back to the client.

5 Actual command output sent to the client.

6 This line displays the return code that is expected when the process has completed.

Example trace of an RSH client using the SEND command

Figure 77 shows a portion of an example of a trace of an RSH client using a SEND command. Short descriptions of numbered items in the trace follow the figure.

```
1
EZA4438I SSCSARAY[0] JOB00043 20 COMPLETED
EZA4385I SSSORT(CTRL): 00000000
EZA4389I SSSORT(init): 00000004
2
EZA4437I SSCSARAY[1] JOB00044 40 WAITING
EZA4438I SSCSARAY[0] JOB00043 20 COMPLETED
EZA4385I SSSORT(CTRL): 00000000
EZA4389I SSSORT(init): 00000004
3
EZA4436I SSCSARAY[1] JOB00044 80 ACTIVE
EZA4438I SSCSARAY[0] JOB00043 20 COMPLETED
EZA4385I SSSORT(CTRL): 00000000
EZA4389I SSSORT(init): 00000004
4
EZA4438I SSCSARAY[1] JOB00044 20 COMPLETED
EZA4385I SSSORT(CTRL): 00000000
5
EZA4392I S99ret: 00000000, A RSHD NEWALTON.RSHD5.JO000044.D0000105.
6
TIME-12:07:02 PM. CPU-00:00:00 SERVICE-1134 SESSION-00:00:00 APRIL 1,2008
EZA4393I S99ret: 00000000
7
EZA4442I SSSORT(next): 00000004 NO MORE OUTPUT ON JES SPOOL
```

Figure 77. Example of a trace of an RSH client using a SEND command

Following are short descriptions of numbered items in the trace:

1 JOB00043 is a previous job that has completed.

2 The 40 indicates that job JOB00044 is waiting for execution. This means that the remote execution server has processed the RSH client request, created a JES job, and has submitted the JOB to JES. The server continues to check the status.

Guidelines: If the status does not change from 40, this could indicate one of the following problems:

- No JES initiator started to process the submitted job class.
- Other jobs might be running in this class that are inhibiting this job from starting.

Make a note of the job number and check JES activity, or check any other jobs that might be running at the same time.

3 The 80 indicates that job JOB00044 is currently active. This means that the remote execution server has checked with JES on the job status and was informed that the job is executing.

Note: If the status does not change from 80 this could indicate that the job is taking too long to run. RSH was not intended to be used for long running jobs. Long running jobs should be submitted using FTP JES processing.

4 The 20 indicates that job JOB00044 is on the output queue. This means that when the server checked with JES on the job status it discovered that the

JES job has completed, and it has been placed on the output queue. The command output should be sent back to the client soon.

Note: See z/OS Communications Server: IP Messages Volume 1 (EZA) for more information about the individual messages in the trace.

5 This line shows the return code from the dynamic allocation of the JES data sent back to the client.

6 Actual command output sent to the client

7 This is the return code expected when there is no more work to do.

Example trace to the JES spool file of the server

Figure 78 on page 597 shows a trace log as it appears on the server. Short descriptions of numbered items in the trace follow the figure.

```
 1  socket 0: EZA4381I Accept 2 from 9.37.217.142 on 512
 2  socket 2: EZA4382I Connecting on 3 to 9.37.217.142:1255
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 0: EZA4381I Accept 4 from 9.37.217.133 on 514
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 4: EZA4382I Connecting on 5 to 9.37.217.133:1584
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 4: EZA4438I SSCSARAY[0] JOB00049 20 COMPLETED
socket 4: EZA4385I SSSORT(CTRL): 00000000
 3  socket 4: EZA4392I S99ret: 00000000, T RSHD USER1.RSHD3.JOB00049.D0000102.?
socket 4: EZA4393I S99ret: 00000000
socket 4: EZA4442I SSSORT(next): 00000004 NO MORE OUTPUT ON JES SPOOL
 4  socket 2: EZA4441I Closing socket 2 from 9.37.217.142 on 512
 5  socket 3: EZA4442I Closing socket 3 to 9.37.217.142:1255
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4436I SSCSARAY[0] JOB00048 80 ACTIVE
socket 2: EZA4438I SSCSARAY[0] JOB00048 20 COMPLETED
socket 2: EZA4385I SSSORT(CTRL): 00000000
socket 2: EZA4392I S99ret: 00000000, T RSHD USER35.RSHD2.JOB00048.D0000102.?
socket 2: EZA4393I S99ret: 00000000
socket 2: EZA4442I SSSORT(next): 00000004 NO MORE OUTPUT ON JES SPOOL
socket 4: EZA4441I Closing socket 4 from 9.37.217.133 on 514
socket 5: EZA4442I Closing socket 5 to 9.37.217.133:1584
```

Figure 78. Example trace to the JES spool file of the server

Following are short descriptions of numbered items in the trace:

Note: Trace messages are preceded by the identification of which socket that trace entry applies too.

The listening socket is socket 0. The EZA4381I message for this socket identifies the socket the request is processed on and the IP address and port that the request was received on. Port 512 is for REXEC requests and port 514 is for RSH requests.

1 Indicates an REXEC request was received on socket 2 for the specified IP address. Subsequent entries beginning with socket 2 indicates activity occurring during the processing of this request. Message EZA441I is issued when this socket is closed.

2 EZA4382I identifies the socket 3 and the IP address and port the server is connecting back to the client on. This port (1255) is in the packet that the client sent to the server. Message EZA4442I is issued when this socket is closed. Common messages are EZA4382I and EZA4442I for this socket.

3 This line shows the return code from the dynamic allocation of the JES data sent back to the client.

4 Indicates the request is completed and the socket is closed.

5 Indicates the error socket is closed

Chapter 22. Diagnosing z/OS UNIX REXEC, RSH, REXECD, and RSHD problems

This topic contains diagnosis information about the z/OS UNIX remote execution protocol (REXEC), remote shell protocol client (RSH), remote execution protocol daemon client (REXECD), and remote shell daemon (RSHD).

Setting up the inetd configuration file

The inetd program is a generic listener program used by such servers as z/OS UNIX TELNETD and z/OS UNIX REXECD. Other servers such as z/OS UNIX FTPD have their own listener program and do not use inetd.

The inetd.conf file is an example of the user's configuration file. It is stored in the /etc directory. Upon startup, the servers for z/OS UNIX TELNETD, rshell, rlogin, and rexec are initiated if they have been defined in /etc/inetd.conf. If it does not include z/OS UNIX TCP/IP applications, add the information shown in Figure 79:

```
#======================================================================
# service | socket | protocol | wait/ | user | server  | server program
# name    | type   |          | nowait|      | program |   arguments
#======================================================================
#
 shell    stream    tcp        nowait  OMVSKERN /usr/sbin/orshd rshd -l
 exec     stream    tcp        nowait  OMVSKERN /usr/sbin/orexecd rexecd -LV
 otelnet  stream    tcp        nowait  OMVSKERN /usr/sbin/otelnetd otelnetd  -LV
 login    stream    tcp        nowait  bpxroot  /bin/rlogind       rlogind   -d  1
 # Add the following line to enable Kerberos for orshd
 kshell   stream    tcp        nowait  OMVSKERN /usr/sbin/orshd orshd -l -k KRB5
```

Figure 79. Adding applications to /etc/inetd.conf

Guideline: For IPv6 support, specify tcp6 for the protocol.

When nowait is specified, the inet daemon issues an accept when a connect request is received on a stream socket. You can specify nowait.max, where max is the maximum number of users allowed to request service in a 60–second interval. The default is 40. If maximum is exceeded, the service's port is shut down. If you expect more than 40 users per minute requesting service, specify the maximum that you expect.

To establish a relationship between the servers defined in the /etc/inetd.conf file and specific port numbers in the z/OS UNIX environment, ensure that statements have been added to ETC.SERVICES for each of these servers. See the sample ETC.SERVICES installed in the /usr/lpp/tcpip/samples/services directory for how to specify ETC.SERVICES statements for these servers.

Guideline: It is important that the service name in /etc/inetd.conf (login in **1**) matches the service name in /etc/services:

```
login 513/tcp
```

The traces for both the z/OS UNIX REXECD server and the z/OS UNIX RSHD server are enabled by options in the inetd configuration file (/etc/inetd.conf). See Figure 80 on page 600.

© Copyright IBM Corp. 2000, 2015

```
#====================================================================
# service | socket | protocol | wait/ | user | server  | server program
# name    | type   |          | nowait|      | program |   arguments
#====================================================================
#
shell      stream    tcp       nowait OMVSKERN /usr/sbin/orshd rshd -d   2
exec       stream    tcp       nowait OMVSKERN /usr/sbin/orexecd rexecd -d   3
```

Figure 80. Setting traces in /etc/inetd.conf

The traces are turned on for both servers by passing a -d argument to the server programs. **2** is the RSHD server and **3** is the REXECD server. All commands executed after the debug flags have been turned on in the inetd configuration file and after the inetd server has reread the file produces trace output.

The trace is written in formatted form to the syslogd facility name daemon with a priority of debug. The trace data can be routed to a file in your Hierarchical File System by specifying the following definition in your syslogd configuration file (/etc/syslogd.conf):

```
  #
  # All ftp, rexecd, rshd
  # debug messages (and above
  # priority messages) go
# to server.debug.a
#
daemon.debug              /tmp/syslogd/server.debug.a
```

In this example, the trace data is written to /tmp/syslogd/daemon.debug.a in your hierarchical file system. See z/OS Communications Server: IP Configuration Reference for more information about syslogd.

For more information about inetd, see z/OS UNIX System Services Planning.

Diagnosing z/OS UNIX REXEC

The following kinds of information can help you diagnose a z/OS UNIX REXEC problem:

- A message beginning with EZYRC
- A code
- An z/OS UNIX REXEC debug trace
- A REXECD debug trace from the foreign host

Activating the z/OS UNIX REXEC debug trace

To activate the z/OS UNIX REXEC debug trace, specify the -d option.

z/OS UNIX REXEC trace example and explanation

The z/OS UNIX REXEC can be invoked using either `rexec` or `orexec`. Enter one of the following commands with either an IP address or a host name.

- IPv4

  ```
  orexec -d -l debfox -p mypwd -s 1512 197.22.190.1 ls -al
  ```

- IPv6

  ```
  orexec -d -l debfox -p mypwd -s 1512 fec0:0:0:12BE::1 ls -al
  ```

The following are examples of the trace output:

- IPv4

```
EZYRC02I Host: 197.22.190.1, user debfox, cmd ls -al, port 1512
```

- IPv6

```
EZYRC02I Host: fec0:0:0:12BE::1, user debfox, cmd ls -al, port 1512
```

```
EZYRC01I  Calling function rexec_af with the following:
EZYRC02I  Host: fec0:0:0:12BE::1, user debfox, cmd ls -al, port 1512
EZYRC19I Data socket = 4, Control socket = 6.
```

EZYRC01I shows that the z/OS UNIX REXEC function has been called in the runtime libraries. EZYRC02I shows the parameters that have been passed to the REXEC() function in the runtime library. EZYRC191 shows the socket descriptor being used for the data connection and the control (or standard error) connection.

Diagnosing z/OS UNIX RSH

The following kinds of information can help you diagnose a z/OS UNIX REXEC problem:

- A code
- A z/OS UNIX RSH debug trace
- An RSHD debug trace from the foreign host

Step for activating the z/OS UNIX RSH debug trace

Activate the z/OS UNIX RSH debug trace.

Procedure

Perform the following step:

Specify the -d option.

Step for invoking z/OS UNIX RSH trace

The z/OS UNIX RSH trace can be invoked using either **rsh** or **orsh**.

Procedure

Enter one of the following commands with either an IP address or a host name:

- IPv4

```
orsh -d -l debfox/mypwd -s 1514 197.22.190.1 date
```

- IPv6

```
orsh -d -l debfox/mypwd -s 1514 fec0:0:0:12BE::1 date
```

Results

The following are examples of the trace output:

- IPv4

```
EZYRC31I Calling function rcmd_af with the following:
EZYRC02I Host: 197.22.190.1, user debfox, cmd date, port 1514
EZYRC19I  Data socket = 4, Control socket = 6.
Thu Apr  3 15:44:11  2003
```

- IPv6

```
EZYRC31I Calling function rcmd_af with the following:
EZYRC02I Host: fec0:0:0:12BE::1, user debfox, cmd date, port 1514
EZYRC19I  Data socket = 4, Control socket = 6.
Thu Apr  3 15:41:11  2003
```

EZYRC31I shows that the local rcmd_af() function has been called. EZYRC02I shows the parameters that have been passed to the rcmd_af() function. EZYRC19I shows the socket descriptor being used for the data connection and the control (or standard error) connection.

Diagnosing z/OS UNIX REXECD

The following kinds of information can help you diagnose a z/OS UNIX REXECD problem:

- A message beginning with EZYRD
- A code
- A z/OS UNIX REXECD debug trace
- A trace from the z/OS UNIX REXECD client

Activating the z/OS UNIX REXECD debug trace

The z/OS UNIX REXEC can be invoked using either `rexecd` or `orexecd`. To activate the z/OS UNIX REXECD debug trace, specify the -d option in the /etc/inetd.conf file.

z/OS UNIX REXECD trace example and explanation

These examples are in the file that is specified in syslogd.conf.

Note: Ensure syslogd is running before collecting these traces and that the file is properly specified.

```
Jun 12 13:31:47  rexecd.851981.: EZYRD31I  MVS OE REXECD BASE
```

The entry is stamped with the date, time, the name of the daemon and the order number of the daemon, the message number (EZAYRD31I), and related information, as shown in the following example.

```
Jun 12 13:31:49  rexecd.851981.: EZYRD03I  Remote address = 9.67.113.61
Jun 12 13:31:49  rexecd.851981.: EZYRD05I  clisecport = 1029
Jun 12 13:31:49  rexecd.851981.: EZYRD08I  User is: user21
Jun 12 13:31:49  rexecd.851981.: EZYRD09I  Command is: ls -l
Jun 12 13:31:49  rexecd.851981.: EZYRD12I  Name is: USER21, user is user21
Jun 12 13:31:49  rexecd.851981.: EZYRD13I  dir is: /u/user21
Jun 12 13:31:49  rexecd.851981.: EZYRD14I  uid is: 21, gid is 0
```

For an explanation of the messages, see z/OS Communications Server: IP Messages Volume 1 (EZA).

Diagnosing z/OS UNIX RSHD

The following kinds of information can help you diagnose a z/OS UNIX RSHD problem:

- A message beginning with EZYRS
- A code
- A z/OS UNIX RSHD debug trace
- A trace from the RSH client

Step for activating the z/OS UNIX RSHD debug trace

The z/OS UNIX RSHD debug trace can be invoked using either `rshd` or `orshd`.

Procedure

Perform the following step to activate the z/OS UNIX RSHD debug trace.

Specify the -d option in the /etc/inetd.conf file.

z/OS UNIX RSHD trace example and explanation

These examples are from the file specified in syslogd.conf.

Restriction: Ensure syslogd is running before collecting these traces and that the file exists and has been properly specified.

```
Jun  9 12:10:04  rshd.4653080.: EZYRS01I  MVS OE RSHD BASE
```

The entry is stamped with the date, time, name of daemon and the order number of the daemon, the message number (EZYRS01I), and related information, as shown in the following example.

```
Jun  9 12:10:06  rshd.4653080.: EZYRS12I  Clisecport = 1020
Jun  9 12:10:06  rshd.4653080.: EZYRS21I  Remote user is: OS2USER
Jun  9 12:10:06  rshd.4653080.: EZYRS22I  Local user is: user21
Jun  9 12:10:06  rshd.4653080.: EZYRS23I  Command is: ls -l
```

For an explanation of the messages, see z/OS Communications Server: IP Messages Volume 3 (EZY).

If the -A option is specified in /etc/inetd.conf, the z/OS UNIX RSHD server does not execute a command when the client host IP address cannot be resolved to a host name.

Resolving garbage errors

There are a few situations where the z/OS UNIX RSHD server might encounter an error so early in the processing of a command that the server has not yet established a proper EBCDIC-to-ASCII translation. In such a situation, the client end user might see garbage data returned to his or her terminal. A packet trace reveals that the response is in fact returned in EBCDIC, which is the reason for the garbage look on an ASCII workstation. This can happen if the z/OS UNIX name resolution has not been configured correctly, so the z/OS UNIX RSHD server, for example, was not able to resolve IP addresses and host names correctly. If your RSH clients encounter such a problem, go back and check your name resolution setup. If you are using a local hosts table, make sure that the syntax of the entries in your hosts file is correct.

Chapter 23. Diagnosing X Window System and Motif problems

This topic describes environment variable XWTRACE that might be useful when diagnosing X Window System and Motif problems. The environment variable, XWTRACE, controls the generation of traces of the socket level communication between Xlib and the X Window System Server.

The following subtopics are included:

- "Trace output when XWTRACE=2"
- "Trace output when XWTRACELC=2" on page 606
- XWTRACE undefined or zero — No trace generated.
- XWTRACE=1 — Error messages.
- XWTRACE>=2 — API function tracing for TRANS functions.

Another environment variable, XWTRACELC, causes a trace of various locale-sensitive routines. If XWTRACELC is defined, a routine flow trace is generated. If XWTRACELC=2, more detailed information is provided.

Guideline: There are no special post-install activities for GDDMXD in z/OS Communications Server. GDDM APAR (PN77391) eliminated these activities for TC/IP Version 3 Release 1. However, if you have an old GDDMXD load library (*tcpip.v3r1*.SEZALNKG) in your LNKLSTxx member in SYSx.PARMLIB, you need to remove that library from the MVS link list, because it is no longer needed.

Following are some examples of X Window System traces.

Trace output when XWTRACE=2

Figure 81 on page 606 shows a typical stream of socket level activity that is generated when an X application running on z/OS UNIX MVS exchanges information with an X Server.

© Copyright IBM Corp. 2000, 2015

```
TRANS(OpenCOTSClient) (/9.2.104.56:0)
TRANS(Open) (1,/9.2.104.56:0)
TRANS(SocketOpenCOTSClient) (inet,9.2.104.56,0)
TRANS(Connect) (3,/9.2.104.56:0)
TRANS(SocketINETConnect) (3,9.2.104.56,0)
TRANS(GetPeerAddr) (3)
TRANS(ConvertAddress)(2,16,7f9d288)
TRANS(SetOption) (3,2,1)
TRANS(SocketWritev) (3,225bc,1)
TRANS(SetOption) (3,1,1)
TRANS(SocketRead) (3,22344,8)
TRANS(SocketRead) (3,22344,8)
TRANS(SocketRead) (3,7f9d368,224)
TRANS(SocketWrite) (3,7f9eb88,60)
TRANS(SocketRead) (3,2249c,32)
TRANS(SocketRead) (3,2249c,32)
TRANS(SocketRead) (3,2249c,32)
TRANS(SocketWrite) (3,7f9eb88,56)
TRANS(SocketRead) (3,22518,32)
TRANS(SocketRead) (3,22518,32)
TRANS(SocketWrite) (3,7f9eb88,80)
TRANS(SocketWrite) (3,7f9eb88,20)
TRANS(SocketRead) (3,223e8,32)
TRANS(SocketRead) (3,223e8,32)
TRANS(SocketDisconnect) (7f9d2c0,3)
TRANS(Close) (3)
TRANS(SocketINETClose) (7f9d2c0,3)
```

Figure 81. Example of X Application trace output when XWTRACE=2

Each line of the trace provides:

- The name of the function involved from x11trans
- Values of the parameters passed to the function

Trace output when XWTRACELC=2

Figure 82 on page 607 is a partial trace showing typical types of information displayed by locale-sensitive routines.

```
lcPubWrap:_XlcCreateLC(C)
 lcCT:_XlcAddCT(ISO8859-1:GL,"(B)
 lcCT:_XlcParseCT
 lcCT:_XlcGetCharSetFromEncoding( (B)
 lcCT:_XlcParseCT returning: 28 charset 0
 lcCharSet:_XlcCreateDefaultCharSet(ISO8859-1:GL,""a)
 lcCT:_XlcParseCharSet
 lcCT:_XlcParseCT
 lcCT:_XlcGetCharSetFromEncoding( (B)
 lcCT:_XlcParseCT returning: 28 charset 0
 lcCharSet:_XlcAddCharSet(ISO8859-1:GL)
 lcCharSet:_XlcGetCharSet(ISO8859-1:GL)
    returned NULL
 lcCT:_XlcAddCT   returning: 7f994d8
:
 (trace statements in this section have been deleted)
 lcCT:_XlcAddCT(CNS11643.1986-1:GL,"$(H)
 lcCT:_XlcParseCT
 lcCT:_XlcGetCharSetFromEncoding( $(H)
 lcCT:_XlcParseCT returning: 2428 charset 0
 lcCharSet:_XlcCreateDefaultCharSet(CNS11643.1986-1:GL,"""+)
 lcCT:_XlcParseCharSet
 lcCT:_XlcParseCT
lcCT:_XlcGetCharSetFromEncoding( $(H)
 lcCT:_XlcParseCT returning: 2428 charset 0
 lcCharSet:_XlcAddCharSet(CNS11643.1986-1:GL)
 lcCharSet:_XlcGetCharSet(CNS11643.1986-1:GL)
    returned NULL
 lcCT:_XlcAddCT   returning: 7f9c4e0
 lcCT:_XlcAddCT(TIS620.2533-1:GR,"-T)
 lcCT:_XlcParseCT
 lcCT:_XlcParseCT returning: 80 charset 0
 lcFile:_XlcResolveLocaleName(C,"","" ""},"2h",)
 lcFile:_XlcResolveName(C,/usr/lib/X11/locale/locale.alias)
 lcFile:_XlcFileName(7f99420,locale)
 lcFile:_XlcResolveLocaleName(C,,"","","")
 lcFile:_XlcResolveName(C,/usr/lib/X11/locale/locale.alias)
 lcFile:_XlcResolveName(C,/usr/lib/X11/locale/locale.dir)
 lcDB:CreateDatabase(/usr/lib/X11/locale/C/XLC_LOCALE)
   0: XLC_XLOCALE, cs0.ct_encoding,      1: ISO8859-1: GL,
   1: XLC_XLOCALE, cs0.wc_encoding,      1: \x00000000,
   2: XLC_XLOCALE, cs0.length,   1: 1,
   3: XLC_XLOCALE, cs0.side,     1: GL:Default,
   4: XLC_XLOCALE, wc_shift_bits,        1: 8,
   5: XLC_XLOCALE, wc_encoding_mask,     1: \x00008080,
   6: XLC_XLOCALE, state_depend_encoding,        1: False,
   7: XLC_XLOCALE, mb_cur_max,   1: 1,
   8: XLC_XLOCALE, encoding_name,        1: STRING,
   9: XLC_FONTSET, fs0.font,     1: ISO8859-1:GL,
  10: XLC_FONTSET, fs0.charset,          1: ISO8859-1:GL,
 ***
 ***
   lcDB:_XlcGetResource(7f99420,XLC_XLOCALE,mb_cur_max)
 lcDB:_XlcGetLocaleDataBase(7f99420,XLC_XLOCALE,mb_cur_max)
 lcDB:_XlcGetResource(7f99420,XLC_XLOCALE,state_dependent)
 lcDB:_XlcGetLocaleDataBase(7f99420,XLC_XLOCALE,state_dependent)
    returning NULL
 lcDB:_XlcGetResource(7f99420,XLC_XLOCALE,encoding_name)
 lcDB:_XlcGetLocaleDataBase(7f99420,XLC_XLOCALE,encoding_name)
    returning lcd=7f99420
 lcFile:_XlcResolveLocaleName(C,"",,"",)
 lcFile:_XlcResolveName(C,/usr/lib/X11/locale/locale.alias)
```

Figure 82. Example of X Application trace output when XWTRACELC=2

Each line of trace provides:

- The name of the locale routine.
- The function invoked within that locale routine.
- Where pertinent, charset name or encoding information, or charset name and encoding information.
- If exiting the invoked function, the trace statement indicates that the function is returning.

Chapter 24. Diagnosing Simple Network Management Protocol (SNMP) problems

This topic explains SNMP-related concepts and terms, including information about how to diagnose SNMP problems, and contains the following subtopics:

- "Overview"
- "Definitions" on page 611
- "Diagnosing SNMP problems" on page 613
- "SNMP traces" on page 633

Overview

The SNMP protocol provides a standardized interface, through which a program on one host (running an SNMP manager) can monitor the resources of another host (running an SNMP agent).

Management information base (MIB)

The information maintained at each agent is defined by a set of variables known as the management information base, or MIB. In addition to the architected list of variables that must be supported by each SNMP agent, an SNMP agent can also support user-defined variables. These user-defined variables that are not part of the architected MIB are known as enterprise-specific MIB variables.

On z/OS Communications Server, the majority of the MIB variables are maintained outside the SNMP agent address space by programs known as SNMP subagents. The subagent program for the TCP/IP-related MIB variables executes in the TCP/IP address space. The subagent program for OMPROUTE-related MIB variables run as part of OMPROUTE, not as a separate application. The subagent program for SLA-related MIB variables runs as a separate application. The subagent program for TN3270E Telnet server MIB variables executes as a separate subtask in the TN3270E Telnet server address space. For a list of all the MIB objects supported by the agent and subagents shipped as part of z/OS Communications Server, see z/OS Communications Server: IP System Administrator's Commands.

In addition, user-written subagent programs can also exist. All subagent programs, whether provided by z/OS Communications Server or user-written, communicate with the SNMP agent over an architected interface known as the Distributed Protocol Interface, or DPI.

When the SNMP agent receives and authenticates a request, it passes the request to the DPI subagent that has registered as the target of the request. You can see this exchange by enabling DPI tracing within the agent.

PDUs

The SNMP protocol is based on the exchange of protocol data units, or PDUs, between the SNMP manager and the SNMP agent.

SNMP has seven types of PDUs:

© Copyright IBM Corp. 2000, 2015

GetRequest-PDU
Sent from the manager to request information from the agent.

GetNextRequest-PDU
Requests the next variable in the MIB tree.

GetBulkRequest-PDU
Requests the next variable in the MIB tree and can also be used to specify multiple successors.

GetResponse-PDU
Sent from the agent to return information to the manager.

SetRequest-PDU
Sent from the manager to alter information at the agent.

Trap-PDU
Sent from the agent to report network events to the manager. A trap is an unconfirmed notification.

Inform-PDU
Sent from an agent to a manager or from a manager to another manager to report a network event. Attempts to confirm delivery are made for Inform-PDUs, not Trap-PDUs.

Functional components

The following topics provide detailed descriptions of the SNMP functional components.

Managers

A manager is a client application that requests management data. z/OS Communications Server provides two management applications, the z/OS UNIX **snmp** command (the **osnmp** command is a synonym for the **snmp** command), and the NetView SNMP command. The **snmp** command is a management application used from the z/OS UNIX shell to monitor and control network elements. The NetView SNMP command provides the same type of functions from the NetView environment.

The **snmp** command runs in a user address space that is created and removed as **snmp** is issued and completed. The NetView SNMP client requires the following started tasks:

- SNMPIUCV subtask of NetView, which runs in the NetView address space and provides the operator interface to SNMP.
- SNMP query stack address space, which provides the protocol support for the SNMP PDUs.

The SNMPIUCV subtask in the NetView address space and the SNMP query stack address space communicate over an IUCV connection.

Agents

An agent is the server that responds to requests from managers. The agent maintains the MIB. z/OS Communications Server supports a trilingual SNMP agent which can understand SNMPv1, SNMPv2c, and SNMPv3 versions of the SNMP protocol. It also communicates with the subagents using DPI1.1 and DPI2.0 protocols.

Subagents

Subagents help the agent by managing a part of the MIB. z/OS Communications Server supports the following subagents:

- TCP/IP subagent that manages TCP/IP-related standard MIB objects and several enterprise-specific MIBs
- OMPROUTE subagent that manages the ospf MIB
- Network SLAPM2 subagent that manages the Network SLAPM2 MIB
- TN3270E Telnet subagent that manages the Enterprise-specific TN3270 Telnet MIB

These subagents communicate with the SNMP agent by using the DPI 2.0 protocol.

Trap forwarder daemon

The trap forwarder daemon on z/OS Communications Server listens for SNMP traps on a specified port and forwards them to other configured ports. This eliminates the port contention problem when multiple managers want to receive notifications at the same well-known port (162) at the same IP address.

Definitions

The SNMP agent, subagents and clients must be configured to TCP/IP before use. If the NetView SNMP client is used, NetView configuration is also required.

Though the SNMP Agent can be started with no configuration files, to implement settings or security other than the defaults, several configuration data sets are required. Most of the configuration data can be configured in several places. Details on the syntax of the statements in the files and the search orders for the files are in z/OS Communications Server: IP Configuration Reference. In the text that follows, uppercase file names (such as OSNMP.CONF) indicate the generic name for the file, which can be any of the places in the search order for the file.

TCP/IP configuration files for SNMP are summarized below. For use of the NetView SNMP command, changes are required for the NetView start procedure and the DSIDMN and DSICMD NCCFLST members of the NetView DSIPARM data set. For additional information, see z/OS Communications Server: IP Configuration Guide.

z/OS UNIX `snmp` command

To use the **`snmp`** command, the following files might be needed:

OSNMP.CONF
: Defines configuration data for sending SNMPv1, SNMPv2, and SNMPv3 requests to SNMP agents. You can name this file as either a z/OS UNIX file system file or an MVS data set (partitioned or sequential).

MIBS.DATA
: Defines textual names for user variables not included in the compiled MIB shipped with the product. You can name this file as either a z/OS UNIX file system file or an MVS data set (partitioned or sequential).

SNMP agent

The SNMP agent (snmpd) uses the following configuration data sets:

OSNMPD.DATA
: Defines initial settings for some MIB variables supported by the agent.

PW.SRC
Defines community names, if the SNMPD.CONF file is not being used. Community name is a mixed-case, case-sensitive field.

SNMPD.BOOTS
Defines SNMPv3 initialization parameters to the SNMP agent if SNMPv3 security is used.

SNMPD.CONF
Defines security configurations and trap destinations to the SNMP agent. Required if SNMPv3 security is used. Can also be used for community-based (SNMPv1 and SNMPv2c) security.

SNMPTRAP.DEST
Defines trap destinations, if the SNMPD.CONF file is not being used.

With z/OS Communications Server, the SNMP agent allows the use of user-based security (SNMPv3) in addition to, or instead of, community-based security (SNMPv1 and SNMPv2c).

For more detailed information about the configuration data sets, see Deciding on SNMP security needs in z/OS Communications Server: IP Configuration Guide.

TCP/IP subagent

The TCP/IP subagent is controlled by statements in the TCP/IP profile. The following statements are important:

SACONFIG
Defines configuration parameters for the TCP/IP subagent.

ITRACE
Specifies the level of tracing used by the TCP/IP subagent.

OMPROUTE subagent

The SNMP OMPROUTE subagent is controlled by statements in the OMPROUTE configuration file. The following statements are important:

ROUTESA_CONFIG
Defines configuration parameters for the OMPROUTE subagent. You can also use the MODIFY ROUTESA command.

OMPROUTE start option -s *n*
Specifies the level of tracing used by the OMPROUTE subagent. You can also use the MODIFY SADEBUG command.

OSPF_INTERFACE
Defines an OSPF interface. The OMPROUTE subagent supports only OSPF MIB (RFC 1850).

Note: At least one OSPF_INTERFACE must be defined.

Network SLAPM2 subagent

The Network SLAPM2 subagent is controlled by start options specified when the subagent is started. The following options are important:

NSLAPM2 start option -c *community*
Defines the community name to be used in connecting to the SNMP agent.

NSLAPM2 start option -P *port*
Defines the port to be used in connecting to the SNMP agent.

NSLAPM2 start option -d *n*
Specifies the level of tracing used by the Network SLAPM2 subagent. You can also use the MODIFY DEBUG,LEVEL command.

TN3270E Telnet subagent

The TN3270E Telnet subagent is controlled by the TNSACONFIG Profile statement. See z/OS Communications Server: IP Configuration Reference for a detailed description of this statement.

SNMP socket call settings

Finally, SNMP makes socket calls that require correct settings in the TCPIP.DATA file. Statements used by SNMP include:

DATASETPREFIX
Can be used in determining the high-level qualifier for agent configuration data sets.

TCPIPJOBNAME
Determines the TCP/IP instance in which SNMP attempts to establish its relationship through the SETIBMOPT socket call. For more information about TCPIPJOBNAME, see the z/OS Communications Server: IP Configuration Reference.

Trap forwarder daemon

The trap forwarder daemon is controlled by the TRAPFWD.CONF file. TRAPFWD.CONF defines the configuration data to forward trap datagrams received on a port to other management applications listening on different ports.

Diagnosing SNMP problems

Problems with SNMP are generally reported under one of the following categories:

- "Abends" on page 614
- Connection problems
  - "Problems connecting the SNMP agent to the TCP/IP address space" on page 614
  - "Problems connecting SNMP agents to multiple TCP/IP stacks" on page 615
  - "Problems connecting subagents to the SNMP agent" on page 615
  - "Problems connecting to the SNMPIUCV subtask" on page 620
  - "Problems connecting the SNMP query stack to the TCP/IP address space" on page 621
- Incorrect output
  - "Unknown variable" on page 621
  - "Variable format incorrect" on page 624
  - "Variable value incorrect" on page 626
- "No response from the SNMP agent" on page 627
- "Report received from SNMP agent" on page 630
- "0.0.0.0 address in traps from the SNMP agent" on page 631
- "I/O error for SNMP PING" on page 631
- "Traps not forwarded by trap forwarder daemon" on page 632

- "Incorrect address in forwarded trap" on page 632

Note: A nonzero return code from the SNMP agent indicates an abnormal termination. For more information, use the SNMP agent traces sent to the SYSLOGD output.

Use the information provided in the following sections for problem determination and diagnosis of errors reported against SNMP.

For additional information, see z/OS Communications Server: IP Configuration Guide and z/OS Communications Server: IP Configuration Reference.

Abends

An abend during SNMP processing should result in messages and error-related information that is sent to the system console. A dump of the error is needed unless the symptoms match a known problem.

Documentation

Code a CEEDUMP DD statement in the PROC used to start the SNMP agent to ensure that a useful dump is obtained in the event of an abend.

Analysis

See z/OS Problem Management or Chapter 3, "Diagnosing abends, loops, and hangs," on page 21, for information about debugging dumps produced during SNMP processing.

SNMP connection problems

This topic describes how to diagnose and correct SNMP connection problems.

Problems connecting the SNMP agent to the TCP/IP address space

Problems connecting the SNMP agent to the TCP/IP address space are indicated by an error message in the agent traces in the syslog daemon output, indicating a socket error. For more information about reading the syslogd traces, see the z/OS Communications Server: IP Configuration Guide.

Documentation:
The following documentation should be available for initial diagnosis of problems connecting the SNMP agent to the TCP/IP address space:

- PROFILE.TCPIP information
- SNMP agent tracing (at level 255) to the syslog daemon output
- TCPIP.DATA information
- OMVS console output for any command responses and traces

Analysis:
Use the following checklist to check for problems connecting the SNMP client or agent address space to the TCP/IP address space:

- Are you connected to the correct TCP/IP address space? This is obviously a concern when running multiple stacks. See "Problems connecting SNMP agents to multiple TCP/IP stacks" on page 615.

  If you get a message "unable to connect to TCPIP JOBNAME," you are not connected to the correct address space. If you have defined two or more stacks, make sure your TCPIPjobname in the TCPIP.DATA data set used by the SNMP

agent matches the NAME field on the SUBFILESYSTYPE statement for ENTRYPOINT(EZBPFIN) in the BPZPRMxx member you used to start z/OS UNIX MVS.

- Did any socket-related errors occur?

  Check the SNMP agent syslogd for socket(), bind(), accept(), or other socket error messages. For example, a bind() failure occurs when one or more of the ports needed by the SNMP agent is already in use. See z/OS Communications Server: IP Configuration Guide for more information about syslogd.
- Is the correct TCPIP.DATA information being used? Is the SYSTCPD DD statement coded in the PROC JCL? Is the RESOLVER_CONFIG environment variable passed on the SNMP agent initialization parameters?

If the problem still occurs after checking the preceding items and making any needed changes, obtain the following documentation for problems connecting the agent.

- Dump of SNMP agent address space.
- Dump of TCP/IP address space.
- The syslogd traces from the agent (using trace level 255). See z/OS Communications Server: IP Configuration Guide for more information about reading the syslogd.

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids for your release of MVS. Obtaining SNMP traces is discussed in "SNMP traces" on page 633.

Problems connecting SNMP agents to multiple TCP/IP stacks

To receive TCP/IP related management data, each TCP/IP stack that is started must run its own SNMP agent. This requires that each agent can find the TCP/IP job name of the TCP/IP stack with which it wants to associate.

Analysis: Use the following checklist to check for problems connecting the SNMP agent to the correct TCP/IP stack:

- Message EZZ6205I indicates that when _iptcpn() was called, it did not return the correct TCPIPjobname for that agent.
  - Check _iptcpn()'s search path.
  - Check to see whether the _BPXK_SETIBMOPT_TRANSPORT environment variable has been set in the cataloged procedure.

  See z/OS Communications Server: IP Configuration Reference for additional information.
- Message EZZ6272I indicates that the setibmopt call failed. This means that _iptcpn() returned a name that z/OS UNIX did not recognize as a PFS. Check the BPXPRMxx member (in SYS1.PARMLIB) used to configure z/OS UNIX.

Problems connecting subagents to the SNMP agent

Problems connecting an SNMP subagent to the SNMP agent are generally indicated by one of the following:

- A socket error at the subagent.
- Authentication failures when the subagent attempts to open a connection.
- A "no such name" response from the SNMP agent when an SNMPv1 manager requests a variable owned by the subagent.
- A "no such object" response from the SNMP agent when an SNMPv2 or SNMPv3 manager requests a variable owned by the subagent.

Documentation:
The following documentation should be available for initial diagnosis of interface connection problems:

- PROFILE.TCPIP information.
- SNMP agent job output, including syslogd output.
- Documentation for the subagent which is not connecting, as follows:
  - TCP/IP subagent syslogd output obtained by specifying the profile statement ITRACE ON SUBAGENT 2 (if the subagent is the TCP/IP subagent).
  - Output of the Netstat HOME/-h command.
  - TCPIP.DATA information.
  - OMPROUTE subagent syslogd output obtained by starting OMPROUTE with the -s1 option or by issuing the MODIFY SADEBUG command to start OMPROUTE subagent tracing (if the subagent is the OMPROUTE subagent).
  - Network SLAPM2 subagent syslogd output obtained by starting the Network SLAPM2 agent with the -d 131 option or by issuing the MODIFY DEBUG,LEVEL command to start Network SLAPM2 subagent tracing (if the subagent is the Network SLAPM2 subagent). The value 131 for -d turns on the following traces.

    1 Trace Network SLAPM2 Subagent Error and System Console Messages

    2 Trace Network SLAPM2 Subagent Warning Message

    128 Trace DPIdebug()level 2
  - TN3270E Telnet subagent syslogd output obtained by specifying the TNSATRACE keyword on the TNSACONFIG profile statement.

Analysis:
Use the following checklist to check for problems connecting an SNMP subagent program to the SNMP agent:

1. Are there multiple TCP/IP stacks active on the same MVS image, are there subagents active for each stack, and are the subagents using UNIX to connect to the agent (as opposed to using TCP)? If so, have you configured a unique UNIX pathname to be used by the subagents connecting to the Agent through UNIX? In a multi-stack environment, each Agent must use a unique UNIX pathname for subagent connections. The default UNIX pathname is /var/dpi_socket. Additional UNIX pathnames can be specified in one of two ways:
   - As the value of the dpiPathNameForUnixStream MIB object in the OSNMPD.DATA configuration file read by the Agent.
   - On the -s start option in the PARM= field of the EXEC JCL statement in the Agent's started procedure.
2. Is the subagent in question the TCP/IP subagent? If so,
   - Is the SACONFIG statement configured correctly?
   - Is SACONFIG disabled?
3. Is the subagent in question the OMPROUTE subagent?
   - Is the OMPROUTE ROUTESA_CONFIG statement configured correctly?
   - Is the OMPROUTE subagent (ROUTESA) disabled?
   - Does the port number match the SNMP agent and OMPROUTE application for the OMPROUTE ROUTESA_CONFIG parameter AGENT=<agent port number>?

- Does the community name (or password) match with the SNMP agent and OMPROUTE application for the OMPROUTE ROUTESA_CONFIG parameter COMMUNITY=<community string>?

4. Is the subagent in question the Network SLAPM2 subagent?
   - Does the port number specified on the -P parameter of the Network SLAPM2 subagent match the port number specified by the SNMP agent?
   - Does the community name (or password) specified on the -c parameter of the Network SLAPM2 subagent match the community name (or password) specified by the SNMP agent?
5. Is the subagent in question the TN3270E Telnet subagent? If so:
   - Is the TNSACONFIG statement configured correctly?
   - Is TNSACONFIG DISABLED specified?
6. If you are using an *hlq*.HOSTS.SITEINFO file (or its z/OS UNIX file system equivalent, /etc/hosts), you must ensure that the IP address in this file for the system on which the agent/subagent are executing matches an interface IP address of the TCP/IP stack to which the agent/subagent are connected. The interface IP addresses for a TCP/IP stack are defined on the HOME profile statement.
7. Is the subagent using the correct IP address to send the connection request to the SNMP agent? The subagent uses the IPv4 primary interface IP address of this stack when sending the connection request to the SNMP agent. The IPv4 primary interface IP address is either the first IP address in the HOME list or the IP address specified on a PRIMARYINTERFACE TCP/IP profile statement. Check the Netstat HOME/-h output to verify the IPv4 primary interface address of the stack. This IP address is the one that is used by the SNMP agent, along with the community name to verify the subagents authority to connect to the SNMP agent.
8. Is the port number correct?
9. Is the community name (or password) correct?

   Guideline: Community name is a mixed-case, case-sensitive field. Many times the client cannot get a response from an agent because the agent has a community string of PUBLIC. Most clients default their community string to `public`.
10. If the SNMP agent is configured for SNMPv3, is the community name configured in the agent SNMPD.CONF file? The subagent can use the community name only if VACM_GROUP, VACM_VIEW, and VACM_ACCESS are defined. For the subagent to connect, the VACM_VIEW must include the dpiPort objects.
11. Did any socket-related errors occur?

    Check the SNMP agent/subagent syslogd for socket(), bind(), accept(), or other socket error messages, particularly error messages related to the DPI connection.
12. If the subagent is using TCP to connect to the SNMP agent then the connection could have been closed by the agent due to a security authorization failure. If the agent security resource name has been defined in the SERVAUTH class, then the subagent must be running on the same TCP/IP stack as the agent and the user ID of the subagent must be permitted to the resource name in order for the connection to succeed. See the SNMP information of the z/OS Communications Server: IP Configuration Guide for a description of the agent security resource name used with TCP connections between SNMP agent and subagent.

If the problem still occurs after checking the preceding items and making any needed changes, obtain the following documentation:

- SNMP agent 255 (trace all) output.
- If the problem is with the TCP/IP subagent, get the subagent traces (level 2). These are turned on by specifying the ITRACE statement in the PROFILE.TCPIP file. This can be done as part of the initial TCP/IP startup. It can also be done after TCP/IP has been started by using the VARY TCPIP command, which is documented in the z/OS Communications Server: IP System Administrator's Commands.
- If the problem is with the OMPROUTE subagent, get the OMPROUTE subagent traces. Turn these on by starting OMPROUTE with the -s1 option or by issuing the MODIFY SADEBUG command to start OMPROUTE subagent tracing.
- If the problem is with the Network SLAPM2 subagent, get the Network SLAPM2 subagent traces. Turn these on by starting the Network SLAPM2 subagent with the -d 131 option or by issuing the MODIFY DEBUG,LEVEL command to start Network SLAPM2 subagent tracing.
- If the problem is with a user-written subagent program, use the DPIdebug() DPI library routine to collect dpi traces in the user-written subagent program. DPIdebug() sends output to the syslogd.
- If the problem is with the TN3270E Telnet subagent, get the subagent traces. These are turned on by specifying the TNSATRACE keyword on the TNSACONFIG statement in the PROFILE.TCPIP file. This can be done as part of the initial TCP/IP startup. It can also be done after TCP/IP has been started by using the VARY TCPIP,,OBEYFILE command, which is documented in the z/OS Communications Server: IP System Administrator's Commands. To enable tracing using the VARY TCPIP,,OBEYFILE command, the subagent must first be disabled and then re-enabled with the TNSATRACE keyword.

The following is a list of things to look for in the SNMP agent trace:

- One of the following incoming SNMP GetRequest-PDU:
  - dpiPortForTcp (1.3.6.1.4.1.2.2.1.1.1) for TCP connect. This is caused by `DPIconnect_to_agent_TCP`
  - dpiPathNameForUnixStream (1.3.6.1.4.1.2.2.1.1.3) for UNIX connect. This is caused by `DPIconnect_to_agent_UNIXstream`.

    Some questions to consider:

    - Was the GetRequest-PDU received? If the GetRequest was not received, was it sent to the right port?

      In the case of the TCP/IP subagent, the value of the AGENT keyword on the SACONFIG statement in the profile must match the value of -p that was specified (or defaulted) when the agent was invoked.
    - Does it have a valid community name in the request?
      - SNMP subagents must use a valid (including correct case) community name as defined in the PW.SRC data set (or SNMPD.CONF data set when using SNMPv3 security) when requesting the dpiPort or dpiPath variable.
      - Note that community name is a mixed-case, case-sensitive field. Specify as follows:
        - For the TCP/IP subagent, specify the community name in the SACONFIG statement.
        - For the OMPROUTE subagent, specify the community name in the ROUTESA_CONFIG statement.

  - For the Network SLAPM2 subagent, specify the community name by way of the -c parameter.
- For the TN3270E Telnet subagent, specify the community name on the TNSACONFIG statement.
- If SNMPv3 is being used, the community name must be defined in the VACM_GROUP statement in the SNMPD.CONF file for the SNMP agent. A VACM_ACCESS statement also needs to be defined to give that group read access to a VACM_VIEW that includes dpiPort objects.
- dpiPathNameForUnixStream defaults to /var/dpi_socket and provides a z/OS UNIX file system pathname used in connecting a DPI subagent with the SNMP agent. The default can be overridden by using the **-s** parameter when starting the agent or by adding an entry for dpiPathNameForUnixStream in the OSNMPD.DATA file.

  A user-written subagent running from a nonprivileged user ID needs write access to the file. Otherwise, a subagent using DPIconnect_to_agent_UNIXstream() would have to be run from an OMVS superuser user ID or other user ID with the appropriate privilege.

- Outgoing GetResponse-PDU for the dpiPort variable:
  - Was the SNMP GetResponse-PDU sent back to the SNMP subagent?
  - Was it sent to the correct IP address?
  - Did it have the correct value for the DPI port?
    - The actual value for the DPI port for TCP can be determined by issuing a Netstat ALL/-A command at the SNMP agent. This displays the port on which the agent is accepting incoming UDP requests.
    - To display the dpiPath name, issue an osnmp get request for dpiPathNameForUnixStream.
- One of the following incoming subagent connections:
  - Message `EZZ6244I Accepted new DPI inet subagent connection on fd fd=xx from inet address xxxx port xxxx.`
  - `EZZ6246I Accepted new DPI inet socket connection on fd=xx`

  Note: *fd=xx* is the number associated with this specific subagent connection. Use it to correlate with later DPI trace messages. The name and number of the port *xxxxx port xxxx*.
- DPI packets transferred for this FD number

The following documentation might also be needed in some cases, but it is suggested that the IBM Software Support Center be contacted before this documentation is obtained:

- Dump of SNMP agent address space
- Dump of TCP/IP address space (for TCP/IP and TN3270E Telnet subagent problems)
- Dump of OMPROUTE address space (for OMPROUTE subagent problems)
- Dump of Network SLAPM2 subagent address space (for Network SLAPM2 subagent problems)
- Dump of user SNMP subagent address space
- Trace from subagent in syslogd

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids. Obtaining SNMP agent traces is discussed in “Starting SNMP agent traces” on page 634.

Problems connecting to the SNMPIUCV subtask

Problems in connecting the SNMPIUCV subtask of Tivoli NetView for z/OS to the SNMP query stack address space are indicated by an error message at the Tivoli NetView for z/OS operator console in response to an SNMP request or an attempt to start the SNMPIUCV subtask.

Documentation:
The following documentation should be available for initial diagnosis of problems connecting the SNMPIUCV subtask to the SNMP query stack:

- PROFILE.TCPIP data set
- SNMP query stack job output, including SYSPRINT output
- NetView log
- SNMPARMS member of DSIPARMS data set

Analysis:
Check for problems connecting the SNMP query stack to the NetView SNMPIUCV subtask:

- Has the SNMP query stack job started successfully?
  - Check the SNMP query stack job output for errors. If the SNMP query stack is started successfully, you should see the message:

    ```
    SQEI001 -- SNMP Query Stack running and awaiting queries...
    ```

    Otherwise, check for errors that might have occurred during socket processing (socket, bind, accept, select, and so on).
- Is the SNMPIUCV subtask started?
  - If not, start the subtask by issuing the command:

    ```
    START TASK=SNMPIUCV
    ```

    from a NetView operator console.
- Was the following message received at the NetView operator console?

  ```
  SNM101W SNMP task (SNMPIUCV) found Query Stack (name) not ready
  ```

  - Is the *name* that the SNMPIUCV subtask is trying to connect to the correct name for the SNMP query stack address space?
    - If not, check the SNMPARMS member of the DSIPARMS data set to make sure that the value specified for the SNMPQE keyword is the correct SNMP query stack address space name.

If the problem still occurs after checking the preceding items and making any needed changes, obtain the following documentation:

- SNMP query stack level 2 trace output
- SNMP query stack IUCV communication trace output

The following documentation might also be needed in some cases, but it is suggested that the IBM Software Support Center be contacted before this documentation is obtained:

- Dump of SNMP query stack address space
- Dump of NetView address space

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids for your release of z/OS. Obtaining SNMP traces is discussed in "SNMP traces" on page 633.

Problems connecting the SNMP query stack to the TCP/IP address space

Problems connecting the SNMP query stack to the TCP/IP address space are usually indicated by an error message in the SNMP client output, indicating either a socket or IUCV error.

Documentation:
The following documentation should be available for initial diagnosis of problems connecting the SNMP query stack to the TCP/IP address space:

- PROFILE.TCPIP data set
- SNMP client output, including SYSPRINT output
- TCPIP.DATA data set

Analysis:
Check the following for problems connecting the SNMP client address space to the TCP/IP address space:

Use the following checklist to check the connection problems:

- Did any socket-related errors occur?

  Check the SNMP query stack job output for socket(), bind(), accept(), or other socket error messages.
- Does job output indicate `RC=1011 received for IUCV_CONNECT to` *`tcpip_name`*?

  Is the *tcpip_name* indicated by the IUCV_CONNECT error the correct name for the TCP/IP address space?
  - Is the correct TCPIP DATA data set being used? (The job output should indicate which data set is being used).
    - Is the SYSTCPD DD statement coded in the PROC JCL?

      Tip: SYSTCPD can be overridden by the global TCPIP.DATA file. See the z/OS Communications Server: IP Configuration Reference for additional information about the search order for the TCPIP.DATA data set.
  - Does the TCPIPJOBNAME keyword in the TCPIP.DATA data set being used have the correct TCP/IP address space name?

If the problem still occurs after checking the preceding items and making any needed changes, obtain SNMP query stack IUCV communication trace output for problems connecting the client.

The following documentation might also be needed in some cases, but it is suggested that the TCP/IP IBM Software Support Center be contacted before this documentation is obtained:

- Dump of SNMP client address space
- Dump of TCP/IP address space

Information about obtaining a dump can be found in the z/OS MVS Diagnosis: Tools and Service Aids manual for your release of z/OS. Obtaining SNMP traces is discussed in "SNMP traces" on page 633.

Incorrect output

Unknown variable

Unknown variable problems are indicated by a **noSuchName** or **noSuchObject** response on an SNMP request. The **noSuchName** response indicates an error

returned on an SNMPv1 request. For SNMPv2 and SNMPv3, more specific errors are returned, such as **noSuchObject** and **noSuchInstance**.

Unknown variable problems are usually caused by one of the following:

- A typographical error in the name or OID
- An incorrect instance number

  Guideline: If the dot-zero (.0) version of this OID contains a non-NULL value, the getnext would return ifNumber.0 and its value. It should be noted that if the dot-zero version of the requested OID is NULL, the getnext returns the first non-NULL value encountered in the MIB tree after ifNumber.0.
- The subagent supporting the MIB object is not started or is not completely connected to the SNMP agent.
- When SNMPv3 is configured, the object is not within the MIB view the user or community can access.

When the NetView SNMP client is used, unknown variable problems are reported when the SNMP client receives either a major error code 2 (internally detected error), minor error code 7 (unknown MIB variable), or a major error code 1 (SNMP agent reported error), minor error code 2 (no such name) in response to an SNMP request.

Documentation:
The following documentation should be available for initial diagnosis of unknown variable problems:

- SNMP syslogd output with traces for both the agent and subagent. For more information about syslogd, see z/OS Communications Server: IP Configuration Guide.
- MIBS.DATA, when snmp is used.
- SNMP query stack job output, when Tivoli NetView for z/OS SNMP is used.
- Tivoli NetView for z/OS log, when Tivoli NetView for z/OS SNMP is used.
- *hlq*.MIBDESC.DATA data set, when Tivoli NetView for z/OS SNMP is used.
- If SNMPv3 security is being used, the SNMP agent configuration file (SNMPD.CONF). If the **snmp** command is the client being used, the **snmp** command configuration file (OSNMP.CONF) might also be needed.
- Include all the configuration files described earlier under "Definitions" on page 611.

Analysis:
Use the following checklist to check for unknown variables at the SNMP agent:

1. Was the variable requested with the correct instance number?

   Variables that are not in a table have an instance number of 0. Variables that are part of a table might have more than one occurrence of the variable value. To get the value of the variable, you need to request a specific instance of the variable. To find the instance number, issue a GET NEXT request; the first occurrence of the variable should be returned.
2. If the variable is not defined in any compiled MIB, is the variable name included in the MIBS.DATA file (for the **snmp** command) or the *hlq*.MIBDESC.DATA file (for the NetView SNMP command)?
3. Did the DPI connection come up successfully?
   a. Check the SNMP agent job output for messages indicating a problem in create_DPI_port.

b. If the DPI port was not successful, no SNMP subagents are able to register MIB variables. The SNMP agent has no knowledge of these unregistered variables and reports them as `noSuchName` for SNMPv1 requests or `noSuchObject` for SNMPv2 and SNMPv3 requests.

4. Has the subagent successfully connected to the SNMP agent?
   a. For subagents shipped as part of z/OS Communications Server, check the MVS operator console for a message indicating that the subagent has completed its initialization.
   b. Issue an **snmp walk** command on the SNMP agent subagent status table. For example, either of the following commands display the subagents that are connected to the z/OS Communications Server SNMP agent and the status of their connections:
      - snmp -v walk saDescription
      - snmp -v walk saStatus

      A value of 1 for saStatus indicates that the subagent connection to the SNMP agent is valid. Following are other possible status values:

```
invalid                 (2)
connecting              (3)
disconnecting           (4)
closedByManager         (5)
closedByAgent           (6)
closedBySubagent        (7)
closedBySubAgentTimeout (8)
closedBySubAgentError   (9)
```

5. If the SNMP agent was configured with SNMPv3 security, is the object within the MIB view of that allowed for that user or community?
   a. Look at the agent SNMPD.CONF file to determine to which VACM_GROUP the user or community name on the failing request belongs. Then examine the VACM_ACCESS statements for that group for the level of security requested on the failing request to determine which MIB views have been permitted to the user or community name.
   b. Alternatively, SNMP agent configuration can be determined from SNMP agent traces if they were set to level 255 at agent initialization.
   c. SNMP agent configuration can also be determined dynamically by issuing snmp walk requests against the agent configuration MIB objects, such as the vacmSecurityToGroupTable and the vacmAcessTable. Reading the values in these tables requires an understanding of how the tables are indexed. See Requests for Comments (RFCs) 2573, 2574, and 2575 for an explanation of the MIB objects containing the SNMP agent configuration.

If the problem still occurs after checking the preceding items and making any needed changes, obtain the following documentation:

For `variable not recognized by manager` messages:

- If the manager is the z/OS UNIX **snmp** command, use the **-d 4** option to obtain all manager traces.
- If the manager is the NetView SNMP command, obtain the SNMP query stack level two output. The SNMP query stack level two trace shows the information flowing between the SNMP query stack and the SNMPIUCV subtask of NetView. Verify in the trace that the variable name being requested is being passed correctly to the SNMP query stack.

For `agent unknown variable`:

- SNMP agent level 15 trace output
- Traces from SNMP subagent programs (if the variable is supported by a z/OS Communications Server subagent)

The SNMP agent level 15 trace shows PDUs between the manager and agent, as well as between the agent and any existing subagents. Look for the following in the trace:

- Is the ASN.1 number received from the manager in the SNMP GetRequest-PDU correct?
- Has a DPI packet registering the requested variable been received?
  1. If not, if you know which subagent program owns the variable, check the subagent program for errors.
  2. If the DPI register has been received, make a note of the FD number for further trace information.
- Were any errors reported for this FD number after the DPI register request was received?
- Was there a DPI information exchange over this FD number as a result of the incoming SNMP GetRequest-PDU?

Another approach to this problem is to look at the agent saMIB variables. This information can be useful when traces are not available. The saMIB variables include the following information:

- An entry for each subagent (including a field for subagent status)
- A table of all trees registered, including:
  - Subagent to which the tree is registered
  - Status of the tree (valid, not valid, and so on)

A description of the saMIB objects can be found in the file samib.mib in the /usr/lpp/tcpip/samples directory.

The following documentation might also be needed in some cases, but it is suggested that the IBM Software Support Center be contacted before this documentation is obtained:

- Dump of SNMP agent address space
- Dump of the subagent responsible for the MIB object whose value is being returned incorrectly.
  - Dump of TCP/IP address space (for TCP/IP and TN3270E Telnet subagent variables)
  - Dump of OMPROUTE address space (for OMPROUTE subagent problems)
  - Dump of Network SLAPM2 subagent address space (for Network SLAPM2 subagent problems)
  - Dump of user subagent address space

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids. Obtaining SNMP traces is discussed later in “SNMP traces” on page 633.

Variable format incorrect

Problems with incorrectly formatted variables are generally reported when the variable value from the GetResponse-PDU is displayed at the manager in the incorrect format (for example, as hexadecimal digits instead of a decimal value or a display string).

Documentation:
The following documentation should be available for initial diagnosis of incorrectly formatted variables:

- MIBS.DATA, when snmp is used
- NetView log, when the NetView SNMP command is used
- *hlq*.MIBDESC.DATA data set, when NetView SNMP is used

Analysis:
Use the following checklist to check incorrect variable format:

1. Is the variable contained in the *hlq*.MIBDESC.DATA data set or MIBS.DATA file?
   a. The SNMP query stack uses the *hlq*.MIBDESC.DATA data set to determine the display syntax of the variable value. NetView SNMP requires that all MIB object names be included in the *hlq*.MIBDESC.DATA data set.
   b. snmp searches the MIBS.DATA file for a MIB name definition. If it is not found, the value in the compiled MIB is used.
2. Is the value listed in the syntax position of the *hlq*.MIBDESC.DATA data set or MIBS.DATA file record for this variable the correct syntax?

   The value specified for syntax (for NetView) is case-sensitive and must be specified in lowercase.
3. For NetView SNMP, is the variable value type specified in message `SNM043I Variable value type:` correct?

   See the z/OS Communications Server: IP System Administrator's Commands section about “Managing TCP/IP Network Resources Using SNMP” for the meanings of the variable value types.

If the problem still occurs after checking the preceding and making any needed changes, obtain the following documentation:

- For the TCP/IP subagent, subagent ITRACE level four output to show that the subagent returned to the SNMP agent.
- For the OMPROUTE subagent, syslogd output obtained by starting OMPROUTE with the -s1 option or by issuing the MODIFY SADEBUG command to start OMPROUTE subagent tracing.
- For the Network SLAPM2 subagent, syslogd output obtained by the Network SLAPM2 subagent with the -d 131 option or the MODIFY DEBUG,LEVEL command to start Network SLAPM2 subagent tracing. .
- For user-written subagents, DPIdebug(2) output, which is sent to the syslogd. For more information about reading the syslogd traces, see the z/OS Communications Server: IP Configuration Guide.
- SNMP query stack level 4 trace output or z/OS UNIX **snmp** command trace level -d 4.
- SNMP manager command output showing the incorrectly formatted variable.
- SNMP agent level 31 trace output shows the DPI packet exchanges between the agent and subagent, as well as the value returned to the manager.
- For the TN3270E Telnet subagent, syslogd output from TNSATRACE keyword on TNSACONFIG profile statement to show what the subagent returned to the SNMP agent.

In the traces, verify that the variable value and syntax are passed correctly in the SNMP GetResponse-PDU from the agent to the SNMP manager.

The following documentation might also be needed in some cases, but it is suggested that the IBM Software Support Center be contacted before this documentation is obtained:

- Dump of the TCP/IP address space (for TCP/IP and TN3270E Telnet subagent problems)
- Dump of SNMP agent address space
- Dump of SNMP query stack address space
- Dump of OMPROUTE address space (for OMPROUTE subagent problems)
- Dump of Network SLAPM2 subagent address space (for Network SLAPM2 subagent problems)

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids. Obtaining SNMP traces is discussed in "SNMP traces" on page 633.

Variable value incorrect

Problems with incorrect variable values are generally reported when the variable value from the GetResponse-PDU that is displayed at the manager contains incorrect information.

Documentation:
The following documentation should be available for initial diagnosis of variables with incorrect values:

- SNMP agent syslogd trace output.
- If the object is supported by the TCP/IP subagent, the syslogd output. Obtain the syslogd output by using the profile statement ITRACE ON SUBAGENT 4.
- MIBS.DATA, if snmp is being used.
- *hlq*.MIBDESC.DATA, if Tivoli NetView for z/OS SNMP is being used.
- Tivoli NetView for z/OS log, if Tivoli NetView for z/OS SNMP is used.

Analysis:
Use the following checklist to check for incorrect variable value:

1. Is the object identifier in the MIB description file correct?
2. Were any errors reported at the SNMP agent when the variable was requested?
3. Is the variable being cached at the SNMP query stack?

   The SNMP query stack uses the *hlq*.MIBDESC.DATA data set to determine the length of time to cache the variable value (or a default time length if the variable is not found in the *hlq*.MIBDESC.DATA data set). If the variable is requested before the caching time is up, the cached value is used instead of obtaining a new value.
4. Is the variable cached at the TCP/IP subagent?

   The TCP/IP subagent caches variable values for the length of time specified by the ibmMvsSubagentCacheTime MIB object, set by default to 30 seconds.
5. Is the variable supported by the Network SLAPM2 subagent? If so, is it being cached? The Network SLAPM2 subagent caches MIB objects for 30 seconds by default, but the cache time can be overridden at subagent initialization time with the -t parameter.
6. Is the variable cached at the TN3270E Telnet subagent? The TN3270E Telnet subagent caches variable values for the length of time specified by the CACHETIME keyword on the TNSACONFIG profile statement, set by default to 30 seconds.

If the problem still occurs after checking the preceding items and making any needed changes, obtain the following documentation:

- SNMP agent level three showing what was returned to the client.
- For the TCP/IP subagent, ITRACE level four trace output showing what the subagent returned to the SNMP agent.
- For the OMPROUTE subagent, syslogd output obtained by starting OMPROUTE with the -s1 option or by issuing the MODIFY SADEBUG command to start OMPROUTE subagent tracing
- For the Network SLAPM2 subagent, syslogd output obtained by the Network SLAPM2 subagent with the -d 131 option or the MODIFY DEBUG,LEVEL command to start Network SLAPM2 subagent tracing.
- For user-written subagents, DPIdebug(2) output which is sent to the syslogd. For more information about reading the syslogd traces, see the z/OS Communications Server: IP Configuration Guide.
- For the TN3270E Telnet subagent, syslogd output from the TNSATRACE keyword on the TNSACONFIG profile statement showing what the subagent returned to the SNMP agent.

In the traces, verify that the variable value is passed correctly from the SNMP subagent to the SNMP agent and from the SNMP agent to the client.

The following documentation might also be needed in some cases, but it is suggested that the IBM Software Support Center be contacted before this documentation is obtained:

- Dump of TCP/IP address space (for TCP/IP and TN3270E Telnet subagent variables).
- Dump of SNMP query stack address space
- Dump of OMPROUTE address space (for OMPROUTE subagent problems)
- Dump of Network SLAPM2 subagent address space (for Network SLAPM2 subagent problems)
- Incorrect values from the TCP/IP subagent are probably due to an error in the TCP/IP stack. In this case, a dump of the TCP/IP address space and a CTRACE from the stack might be useful. You can also use the Netstat command to verify that the TCP/IP subagent is reporting what the TCP/IP stack believes the value to be.

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids. Obtaining SNMP traces is discussed in “SNMP traces” on page 633.

No response from the SNMP agent

Problems receiving a response from the SNMP agent are generally reported when an SNMP request is issued from a manager but no response from the agent is received. This is reported as a timeout message.

Documentation

The following documentation should be available for initial diagnosis when no response is received from the agent:

- If the z/OS UNIX**snmp** command is being used, any command responses and traces at trace level -d 4.
- If the NetView SNMP command is being used, NetView console output or command responses.

- SNMP agent syslogd trace output at trace level 35.
- The OSNMP.CONF file (if the z/OS UNIX**snmp** command is the manager).
- PW.SRC or SNMPD.CONF file being used by the SNMP agent.
- If the SNMP Manager API is being used, any command responses and trace output for level 15 (SNMP_LOG_ALL and SNMP_LOG_INTERNAL).

Analysis

Use the following checklist when no response is received from an agent:

1. Is the SNMP agent running?
2. Is a path to the agent available? Try issuing a PING request to the IP address of the agent.
3. What is the timeout value? For example, the timeout value on the **snmp** command defaults to three seconds. Trying the request again with a larger timeout value, such as 15 seconds, might result in an answer.
4. Does the request use the correct port number and IP address?
5. Were any errors reported at the SNMP agent when the variable was requested?
6. If community-based security is being used, is the correct community name (including correct case) being used in the request?
7. Is the community name defined for the IP address from which the request originates? For example, a community name defined only for IPv4 addresses is not be usable from an IPv6 address.
8. Is the community name defined for the SNMP version of the request? If the PW.SRC file is being used for community name definitions, community names are usable for both SNMPv1 and SNMPv2c requests. If the SNMPD.CONF file is being used for community name definitions, separate definitions are required to allow the use of the community name for both SNMPv1 and SNMPv2c requests. Note that the **snmp** command defaults to sending SNMPv1 requests. To send an SNMPv2c request using the **snmp** command, an entry is required in the OSNMP.CONF file and the **snmp** command must be issued with a -h value that refers to an entry in the OSNMP.CONF file.
9. Does the agent support the SNMP version of the request? The z/OS Communications Server supports SNMPv1, SNMPv2c and, if configured with SNMPD.CONF, SNMPv3.
10. Are you using SNMPv3 with Advanced Encryption Standard (AES) encryption for the privacy protocol? For AES encryption, the z/OS Integrated Cryptographic Service Facility (ICSF) must be active. For detailed information about configuring ICSF, see z/OS Cryptographic Services ICSF Administrator's Guide.

If the problem still occurs after checking the preceding items and making any needed changes, obtain SNMP agent level seven trace output documentation.

Check the following in the SNMP agent traces:

1. Was the SNMP request PDU received by the agent?
2. Did it have a valid community name? Note that community name is case-sensitive and mixed-case.
3. Was the IP address of the manager the expected IP address?
4. Was an SNMP GetResponse-PDU sent back to the manager?
5. Was an AuthenticationFailure trap generated?

Guideline: For these traps to be generated, you must first provide the trap destination information in either the SNMPTRAP.DEST or SNMPD.CONF file. Then, provide OSNMPD.DATA information where the snmpEnableAuthenTraps MIB object is set to 1, to enable the authentication traps. For detailed information about enabling traps, see z/OS Communications Server: IP Configuration Reference.

The following documentation might also be needed in some cases, but contact the IBM Software Support Center before this documentation is obtained:

- Dump of SNMP agent address space
- Dump of the TCP/IP address space (for TCP/IP and TN3270E Telnet subagent problems)
- Dump of OMPROUTE address space (for OMPROUTE subagent problems)
- Dump of Network SLAPM2 subagent address space (for Network SLAPM2 subagent problems)

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids. Obtaining SNMP traces is discussed in "SNMP traces" on page 633.

SNMPv3 AES encryption or decryption failures

SNMP problems might be caused by SNMPv3 AES privacy protocol encryption or decryption failures. The support of SNMPv3 AES privacy protocol requires the services of the z/OS Integrated Cryptographic Service Facility (ICSF). Before you configure the AES privacy protocol, ensure that ICSF is active.

Documentation

The following documentation should be available for initial diagnosis:

- Any command responses and traces at trace level -d 4 if the z/OS UNIX**snmp** command is being used.
- SNMP agent syslogd trace output at trace level 35.
- The OSNMP.CONF file (if the z/OS UNIX **snmp** command is the manager).
- PW.SRC or SNMPD.CONF file that the SNMP agent is using.
- If the SNMP Manager API is being used, any command responses and trace output for level 15 (SNMP_LOG_ALL and SNMP_LOG_INTERNAL).

Analysis

Use the following checklist to determine whether AES encryption or decryption failures are occurring, based on the SNMP function that you use. For failures other than ICSF not being active, contact the IBM Software Support Center with the documentation.

1. Is ICSF active? For the SNMP agent, the z/OS UNIX **snmp** command, and the SNMP manager API, check the syslog trace output or traces for the following entries to see whether ICSF is active. If the return code is 12 and the reason code is 0, ICSF was not active when you attempted to encrypt or decrypt by using the function.
   - `Encrypting AESCFB128 completed return=12 reason=0`
   - `Decrypting AESCFB128 completed return=12 reason=0`
2. SNMP agent. AES encryption or decryption errors in the SNMP agent result in one of the following situations:
   - Timeout messages are issued by the manager because the agent did not respond to a request.

- The manager does not receive notifications.

To determine whether AES encryption or decryption errors are the cause of these problems, look for the following entries in the agent's syslog trace output. One or both might be present.

- `Encrypting AESCFB128 completed return=nn reason=nn`
- `Decrypting AESCFB128 completed return=nn reason=nn`

3. z/OS **snmp** command.
   - AES encryption errors result in the command issuing the following message:
     `EZZ3301I Error return from SnmpSendMsg()`
   - AES decryption errors result in a timeout message.

   To determine whether AES encryption or decryption errors are the cause of these problems, look for the following entries in the trace output of the command. One or both might be present.
   - `Encrypting AESCFB128 completed return=nn reason=nn`
   - `Decrypting AESCFB128 completed return=nn reason=nn`
4. SNMP manager API.
   - AES encryption errors result in the snmpSendRequest failing with a return code of SNMP_MGR_RC_ENCODE_ERROR (-13).
   - AES decryption errors result in the snmpSendRequest failing with a return code of SNMP_MGR_RC_USM_DECRYPTION_ERROR (-68).

   To determine whether AES encryption or decryption errors are the cause of these problems, look for the following entries in the syslog trace output of the manager API. One or both might be present.
   - `Encrypting AESCFB128 completed return=nn reason=nn`
   - `Decrypting AESCFB128 completed return=nn reason=nn`

Report received from SNMP agent

With SNMPv3, certain error conditions detected on a request are sent back from the SNMP agent to the SNMP manager as a report. Some reports are expected as part of normal processing, but most often they indicate an error condition.

For the **snmp** command, some reports occur during normal processing, such as the usmStatsUnknownEngineIDs condition, which occurs as the **snmp** command performs discovery processing to learn the SNMP stackID of the agent with which it is communicating. Normal processing reports are not displayed by snmp unless debug tracing is active. Reports that indicate error conditions are typically displayed using the EZZ33431 message. For example, when an attempt is made to use a USM user with an authentication key that does not match the key that is configured at the SNMP agent, the usmStatsWrongDigests report is received.

Figure 85 on page 638 shows the output received by an SNMP manager when the authentication key sent by an **snmp** command did not match the key defined at the agent. The command issued in the z/OS UNIX shell was:

```
$ snmp -h v374 -v walk usmUserStatus

EZZ33431 Report received : usmStatsWrongDigests
EZZ33011 Error return from SnmpRecvMsg()
```

Following are other common reports:

usmStatsUnknownUserNames
Indicates a request was received for a user that is not defined at the SNMP agent.

usmStatsUnsupportedSecLevels
Indicates a request was received for a defined user, but the user was not configured at the SNMP agent to use the security level specified in the request.

usmStatsDecryptionErrors
Indicates an encrypted request was received at the SNMP agent, but the request could not be decrypted. This can be the result of an invalid privacy key.

0.0.0.0 address in traps from the SNMP agent

SNMPv1 traps contain the IP address of the originating agent encoded as part of the protocol data unit. The address field is four bytes long. If SNMPv1 traps are received from the z/OS Communications Server SNMP agent with an agent address of 0.0.0.0, it is most likely due to the fact that the agent obtained an IPv6 address for itself when it initialized. To avoid this situation, the SNMP agent can be started with the -A parameter to request that the SNMP agent obtain an IPv4 address for itself when initializing. See the z/OS Communications Server: IP Configuration Reference for more information.

I/O error for SNMP PING

Tivoli NetView for z/OS users can issue a PING request by using SNMP PING. SNMP I/O error problems are reported when a major return code of 2 (internally detected error) and a minor return code of 4 (some I/O error occurred) are received when issuing an SNMP PING. This type of problem is generally caused by an error in the PROFILE.TCPIP data set.

Documentation

The PROFILE.TCPIP data set should be available for initial diagnosis of SNMP I/O problems.

Additional documentation that might be needed later is described in "Analysis."

Analysis

Obtain the following documentation:

- SNMP query stack job SYSPRINT output
- SNMP query stack level two trace output
- SNMP query stack IUCV communication trace output

The following documentation might also be needed in some cases, but it is suggested that TCP/IP customer support be contacted before this documentation is obtained:

- Dump of SNMP address space
- TCP/IP packet trace

Information about obtaining a dump can be found in z/OS MVS Diagnosis: Tools and Service Aids manual for your release of MVS. Obtaining SNMP traces is discussed "SNMP traces" on page 633.

Traps not forwarded by trap forwarder daemon

Problems with traps not getting forwarded by the trap forwarder daemon are most likely the result of configuration errors or problems in the network.

Documentation

The following documentation should be available for initial diagnosis:

- TRAPFWD.CONF file
- Trapfwd traces, level 3
- Traces from the sending agent (the originator of the trap)
- Trace from the receiving client (the target of the forwarded trap)

Analysis

Use the following checklist to check for traps not forwarded by trap forwarder daemon:

1. Is the target address correctly configured in the TRAPFWD.CONF file?

   If the target is designated by a host name, check the trapfwd trace to determine whether the hostname was correctly resolved to an IP address. If the target is designated by an IPv6 colon-hexadecimal address, then your TCP/IP stack must be running with IPv6 support. If the stack is not running with IPv6 support, then the trap forwarder daemon cannot forward traps to IPv6 listener addresses.
2. Is the trap being received at the trap forwarder daemon?

   If trapfwd traces indicate the trap is not being received at the trapfwd daemon, examine traces from the SNMP agent sending the trap. Determine whether the SNMP agent did send the trap.
3. Are there network problems between the trap forward daemon and the target client?

   By issuing an SNMP GET request at the target client to the SNMP agent on the same host as the trap forward daemon, you can determine whether UDP packets are correctly reaching the client.
4. Are the UDP packets being discarded due to congestion at the TCP/IP stack?

   If the trapfwd trace indicates that the trap is correctly being sent from the trap forwarder daemon to the target client, but the trap is not being received, consider setting NOUDPQueuelimit on the UDPCONFIG statement. This is used to specify that UDP should not have a queue limit and would prevent traps from being lost due to congestion.

If the above analysis does not correct the problem, the following documentation should be gathered and the IBM Software Support Center should be contacted:

- UDP packet trace on the TCP/IP stacks where the originating SNMP agent, the trap forwarder daemon, and the target client are running.

Incorrect address in forwarded trap

Documentation

The following documentation should be available for initial diagnosis:

- TRAPFWD.CONF file
- Trapfwd traces, level 3
- Traces from the sending agent (the originator of the trap)
- Trace from the receiving client (the target of the forwarded trap)

Analysis

Use the following checklist to check for an incorrect address in forwarded trap:

1. What is the version of the SNMP trap?

   In the case of SNMPv1 traps, the address from which the trap originated is encoded within the trap packet. A manager that needs the originating address should look into the SNMP packet to get the address.

   If the address is 0.0.0.0, the most likely cause is that the trap originated at an IPv6 address. If the trap originated at the z/OS SNMP Agent, see "0.0.0.0 address in traps from the SNMP agent" on page 631.

   In the case of SNMPv2 traps, the originating address is not encoded within the trap PDU. If a manager uses the address from which the trap packet is received, it would not be the originating address but the address at which the trap forwarder daemon is running. If the manager needs the originating address if SNMPv2 traps, the trap forwarder should be configured to append the originating address to the trap and the manager should be capable of reading the address from the end of the received trap packet. For more information about the format in which the address is appended, see the z/OS Communications Server: IP User's Guide and Commands.

2. Is it an SNMPv2 trap?

   Check to see whether the ADD_RECVFROM_INFO is specified correctly in the TRAPFWD.CONF file. If it is not specified, then add the option to the configuration file. Note, the receiving manager must be capable of processing the RECVFROM information at the end of the trap packet.

If the above analysis does not correct the problem, collect the above documentation and contact the IBM Software Support Center.

SNMP traces

There are several types of traces that can be useful in identifying the cause of SNMP problems:

- Manager traces
- Agent traces
- Subagent traces
- TRAPFWD traces

These traces are discussed in the following topics.

Starting manager traces

To obtain traces when the SNMP manager being used is the z/OS UNIX **snmp** command, issue snmp with the -d option. You can specify a trace level of 0-4. A trace level of 0 provides no tracing, while a level 4 provides all traces. Tracing for snmp is done on a per-request basis. Traces return to the console, but they can be redirected to a file issuing the OMVS redirect operand (>).

When NetView SNMP is being used, traces for the SNMP Query Engine can be obtained by starting the SNMP Query Engine and specifying -d trace_level where trace_level is one of the following:

1 Display errors.

2 In addition to trace level 1, also display SNMP query stack protocol

packets exchanged between the SNMP query stack and the SNMPIUCV subtask, with the exception of TRAP packets sent to NetView from the query stack.

3 In addition to trace level 2, also display decoded SNMP protocol packets sent and received along with some additional informational messages.

4 In addition to trace level 3, display the BER-encoded packets received from NetView or from an SNMP agent. Also, add display of SNMP query stack protocol packets for TRAPs sent from the query stack to NetView.

For example:

```
//SNMPQE   EXEC PGM=SQESERV,REGION=4096K,TIME=1440,PARM='-d 3'
```

Also, the -it option can be used to obtain a trace of IUCV communication.

SNMP Query Engine trace output is sent to the SYSPRINT DD specified in the Query Engine JCL.

Starting SNMP agent traces

If agent is not running

If the SNMP agent is not already running, specify the -d parameter when you invoke the agent. You specify this parameter based on the method by which you start the SNMP agent:

- Using the start options in the JCL used to start the SNMP agent (more common). For example,

  ```
  //OSNMPD    EXEC PGM=EZASNMPD,REGION=4096K,TIME=1440,PARM='-d 8'
  ```

- Using the z/OS UNIX shell, using the **snmpd** command. For example:

  ```
  snmpd -d 255 &
  ```

Use one of the following trace levels or a combination of them:

1 Trace SNMP requests

2 Trace SNMP responses

4 Trace SNMP traps

8 Trace DPI packets level 1

16 Trace DPI internals level 2

32 Internal trace (debug level 1)

64 Extended trace (debug level 2)

128 Trace DPI internals level 2

Combining trace levels: To combine trace levels, add trace level numbers. For example, to request SNMP requests (level 1) and SNMP responses (level 2), you would request trace level 3.

Trace records are sent to the file specified by the daemon.debug entry in the SYSLOG configuration file. For more information see z/OS Communications Server: IP Configuration Guide.

If agent is already running

You can use the MVS MODIFY command to start and stop trace dynamically. Use of this support is restricted to the users with MODIFY command privilege.

If you start the agent from JCL, you have no difficulty knowing the procname. However, if you start the agent from the z/OS UNIX shell, the agent generates a message to syslogd. This message indicates the job name the agent is running; this is the job name to specify on the MODIFY command.

For example, assume the procname is OSNMPD and you want to change the trace level to 3 (tracing SNMP requests and responses). Enter:

```
MODIFY OSNMPD,trace,level=3
```

For more information about using the MVS MODIFY command, see z/OS Communications Server: IP System Administrator's Commands.

Starting TCP/IP subagent traces

To start the TCP/IP subagent traces, code the ITRACE statement in the PROFILE.TCPIP data set or in the data set specified on the VARY TCPIP,,OBEYFILE command. For more information, see z/OS Communications Server: IP Configuration Reference.

```
ITRACE ON SUBAGENT level
```

where *level* is one of the following values:

1 General subagent trace information.

2 General subagent trace information plus DPI traces.

3 General subagent trace information plus extended dump trace. This level provides storage dumps of useful information, such as storage returned by the IOCTL calls.

4 General subagent trace information, plus extended dump trace and DPI traces.

The trace output is sent to the syslogd. Trace records are found in the file specified by the daemon.debug entry in the SYSLOG configuration file. For more information, see z/OS Communications Server: IP Configuration Guide.

To stop TCP/IP subagent traces, code the ITRACE statement in the PROFILE.TCPIP data set or in the data set specified on the VARY TCPIP,,OBEYFILE command:

```
ITRACE OFF SUBAGENT
```

For more information about the VARY command, see z/OS Communications Server: IP System Administrator's Commands.

Starting OMPROUTE subagent traces

To start OMPROUTE subagent tracing, start OMPROUTE with the -s1 option or issue the MODIFY SADEBUG command. Output is sent to syslogd. For details, see “Starting OMPROUTE tracing and debugging from the z/OS UNIX System Services shell” on page 775 and “Starting OMPROUTE tracing and debugging using the MODIFY command” on page 776.

Starting Network SLAPM2 subagent traces

To start Network SLAPM2 subagent tracing, start the Network SLAPM2 subagent with the -d option or by issuing the MODIFY DEBUG,LEVEL command. Output is sent to syslogd.

The Network SLAPM2 subagent trace levels are 0-511. There are nine levels of tracing provided. Each level selected has a corresponding number. The sum of the numbers associated with each level of tracing selected is the value which should be specified as level. After the Network SLAPM2 Subagent is started, tracing options can be dynamically changed using the MVS MODIFY command.

The numbers for the trace levels are:

0 No tracing

1 Trace Network SLAPM2 Subagent Error and System Console Messages

2 Trace Network SLAPM2 Subagent Warning Messages

4 Trace Network SLAPM2 Subagent Informational Messages

8 Trace Network SLAPM2 Subagent Internal statistics table

16 Trace Network SLAPM2 Subagent Internal monitor table

32 Trace Network SLAPM2 Subagent Internal traps

64 Trace Network SLAPM2 Subagent Internal monitoring

128 Trace Network SLAPM2 Subagent Internal Policy Agent API

256 Trace DPIdebug()level 2

Starting TN3270E Telnet subagent traces

To start the TN3270E Telnet subagent traces, code the TNSATRACE keyword on the TNSACONFIG statement in the PROFILE.TCPIP data set or in the data set specified on the VARY TCPIP,,OBEYFILE command. For more information, see z/OS Communications Server: IP Configuration Reference.

If the subagent is not currently tracing, the subagent must first be disabled. Disable the subagent by using the VARY TCPIP,,OBEYFILE command where the data set for the command contains:

```
TELNETGLOBALS
  TNSACONFIG DISABLED
ENDTELNETGLOBALS
```

Then re-enable the subagent by using the VARY TCPIP,,OBEYFILE command where the data set for the command contains:

```
TELNETGLOBALS
  TNSACONFIG ENABLED TNSATRACE
ENDTELNETGLOBALS
```

The trace output is sent to the syslogd. Trace records are found in the file specified by the daemon.debug entry in the SYSLOG configuration file. For more information, see z/OS Communications Server: IP Configuration Guide.

Starting TRAPFWD traces

The following topics provide information about starting TRAPFWD traces.

If TRAPFWD is not running

If TRAPFWD is not already running, specify the -d parameter during startup. You can start the TRAPFWD trace in one of the following ways:

- Through the start options in the JCL used to start the TRAPFWD. For example,

  ```
  //TRAPFWD EXEC PGM=EZASNTRA,REGION=4096K,TIME=NOLIMIT,
  //PARM='POSIX(ON) ALL31(ON)/-d 3'
  ```

- Through OMVS, using the **trapfwd** command. For example,
  ```
  trapfwd -d 3 &
  ```

Use one of the following trace levels:

1 Minimal tracing, trace address from which the trap is received.

2 In addition to 1, trace addresses to which the trap packet is forwarded.

3 In addition to 2, trace trap packets.

Trace records are sent to the file specified by the daemon.debug entry in the SYSLOG configuration file. For more information see z/OS Communications Server: IP Configuration Guide.

If TRAPFWD is already running

You can use the MVS MODIFY command to start and stop the trace dynamically. Use of this support is restricted to users with MODIFY command privilege.

If you start the trapfwd from JCL, you have no difficulty knowing the procname. However, if you start the trapfwd from OMVS, the trapfwd generates a message to syslogd. This message indicates the job name the trapfwd is running; this is the job name to specify on the MODIFY command.

For example, assume that the procname is TRAPFWD and you want to change the trace level to 3. You would enter the following:

```
MODIFY TRAPFWD,trace,level=3
```

For more information about using the MVS MODIFY command, see z/OS Communications Server: IP System Administrator's Commands.

Trace examples and explanations

The following examples are shown in this topic:

- Agent trace
- TCP/IP subagent traces
- TRAPFWD trace
- NetView SNMP Query Engine trace
- NetView SNMP Query Engine IUCV Communication trace

Figure 83 on page 638 was produced by using **snmp get sysUpTime.0**. When the SNMP agent is tracing responses, it makes the following entry in the syslogd output file:

```
Dec 19 15:55:38  snmpagent.9.: SNMP logging data follows ==============

Dec 19 15:55:39  snmpagent.9.: Log_type:        snmpLOGresponse_out
Dec 19 15:55:39  snmpagent.9.:    send rc:      0
Dec 19 15:55:39  snmpagent.9.:    destination:  UDP 127.0.0.1 port 5000
Dec 19 15:55:39  snmpagent.9.: version:         SNMPv1
Dec 19 15:55:39  snmpagent.9.: community:       public
Dec 19 15:55:39  snmpagent.9.:  ('70 75 62 6c 69 63'h)
Dec 19 15:55:39  snmpagent.9.: addressInfo:     UDP 127.0.0.1 port 5000
Dec 19 15:55:39  snmpagent.9.: PDUtype:         GetResponse ('a2'h)
Dec 19 15:55:39  snmpagent.9.: request:         1
Dec 19 15:55:39  snmpagent.9.: error-status:    noError (0)
Dec 19 15:55:39  snmpagent.9.: error-index:     0
Dec 19 15:55:39  snmpagent.9.: varBind oid:
Dec 19 15:55:39  snmpagent.9.: OBJECT_IDENTIFIER
Dec 19 15:55:39  snmpagent.9.: 1.3.6.1.2.1.1.3.0
Dec 19 15:55:39  snmpagent.9.:         name:    sysUpTime.0
Dec 19 15:55:39  snmpagent.9.:         value:
Dec 19 15:55:39  snmpagent.9.: TimeTicks
Dec 19 15:55:39  snmpagent.9.: 5900 - 59.00 seconds
Dec 19 15:55:39  snmpagent.9.: End of SNMP logging data:
```

Figure 83. SNMP agent response trace

In the following scenario, the SNMP agent attempted to initialize, but it was not successful. The port it was using was already in use. The trace shown in Figure 84 was obtained with SNMP agent tracing set to 7.

```
Dec 19 11:57:52  snmpagent.16777227.: EZZ6235I socket function failed for
SNMP inet udp socket; EDC5112I Resource temporarily unavailable.
Dec 19 11:57:52  snmpagent.16777227.:  ... errno = 112, errno2 =12fc0296
```

Figure 84. SNMP agent trace of unsuccessful initialization

Note: Errno 112 translates to `Resource temporarily unavailable.` The errno is used primarily by IBM service in diagnosing the error. In this case, issue the Netstat CONN/-c command to determine whether TCP/IP is running and, if so, which ports are in use.

Figure 85 shows the trace produced for the agent when the authentication key sent by a manager does not match the key defined at the agent. The command receives a report indicating usmStatsWrongDigests.

```
IDSTMVS.S@AU1104.SOURCE.S@AGV123(1624): rc=-65 (SNMP_RC_USM_WRONG_DIGEST)
   from snmp_process_message()
```

Figure 85. SNMP messages and agent trace for nonmatching key

Figure 86 on page 639 shows the output received by an SNMP manager and the trace produced for the agent when the operator attempted to retrieve data not within the defined view. The command issued in the z/OS UNIX shell was:

```
snmp -h v374a -v get usmUserStatus.12.0.0.0.2.0.0.0.0.9.67.35.37.2.117.49
```

```
OUTPUT RECEIVED BY THE MANAGER
usmUserStatus.12.0.0.0.2.0.0.0.0.9.67.35.37.2.117.49 = noSuchObject

AGENT TRACE
IDSTMVS.S@AU1104.SOURCE.S@AGV123(1624): RC=-30 (SNMP_RC_NOT_IN_VIEW)
from snmp_process_message()
```

Figure 86. SNMP messages and agent trace when data not in defined view

The following return codes in SNMP agent traces typically indicate configuration errors:

- SNMP_RC_NOT_AUTHENTICATED - indicates the SNMP agent received an SNMPv1 or SNMPv2c request with a community name that was not valid for use by the IP address making the request.
- SNMP_RC_NOT_IN_VIEW - indicates the SNMP agent received an SNMPv3 request for a MIB object that is not defined to be accessible by the community name or user name making the request.
- SNMP_RC_USM_UNKNOWN_USERNAME - indicates the SNMP agent received an SNMPv3 request for a username not configured at the SNMP agent.
- SNMP_RC_USM_WRONG_DIGEST - indicates the SNMP agent received an SNMPv3 request for which the authentication key for the user making the request was not valid.
- SNMP_RC_USM_DECRYPTION_ERROR - indicates the SNMP agent received an encrypted request, but the request could not be decrypted because the encryption key for the user making the request was not valid.
- SNMP_RC_USM_UNSUPPORTED_SECLEVEL - indicates the SNMP agent received a request for a defined user, but the user was not configured to use the security level specified in the request.

SNMP agent traces

TCP/IP subagent trace

When requests for MIB variables supported by the TCP/IP subagent fail with an indication that the variable is not supported (noSuchName or noSuchObject), one possibility is that the TCP/IP subagent was unable to connect to the SNMP agent.

Figure 87 on page 640 illustrates a scenario where the subagent is unable to connect because the password it is using is not accepted by the SNMP agent. (The password used by the subagent is defined or defaulted on the SACONFIG statement in the TCP/IP profile.) The following traces were obtained with SNMP agent traces set to 15 and the subagent traces (as set on the ITRACE profile statement) set to 3.

```
Apr  4 16:28:17 MVS097 snmpagent[67108869]: EZZ6225I SNMP agent: Initialization complete
Apr  4 16:28:21 MVS097 M2SubA[50331651]: VS.2575 do_connect_and_open: DPIconnect_to_agent_UNIXstream rc -2.
Apr  4 16:28:21 MVS097 M2SubA[50331651]: VS.1320 do_open_and_register: Connect to SNMP agent failed, will
keep trying
Apr  4 16:28:21 MVS097 M2SubA[50331651]: VS.1340 do_open_and_register: issue selectex, interval = 10 seconds
Apr  4 16:28:31 MVS097 M2SubA[50331651]: VS.2543 do_connect_and_open .... getting agent info from config
Apr  4 16:28:31 MVS097 M2SubA[50331651]: 08B7A4A0  82818497  A6404040  40404040  40404040   *badpw        *
Apr  4 16:28:31 MVS097 M2SubA[50331651]: 08B7A4B0  40404040  40404040  40404040  40404040   *             *
Apr  4 16:28:31 MVS097 M2SubA[50331651]: 08B7A49C  000000A1                                  *....         *
Apr  4 16:28:31 MVS097 M2SubA[50331651]: VS.2556 do_connect_and_open: port 161
Apr  4 16:28:31 MVS097 M2SubA[50331651]: VS.2567 do_connect_and open: hostname_p => 9.67.35.37
Apr  4 16:28:31 MVS097 snmpagent[67108869]: IDSTMVS.SO064350.SOURCE.S@AGV123(1623): rc=-14
(SNMP_RC_NOT_AUTHENTICATED) from snmp_process_message()
Apr  4 16:28:34 MVS097 snmpagent[67108869]: IDSTMVS.SO064350.SOURCE.S@AGV123(1623): rc=-14
(SNMP_RC_NOT_AUTHENTICATED) from snmp_process_message()
```

Figure 87. SNMP subagent trace

NetView SNMP query engine trace

This topic discusses the output produced by the SNMP query stack trace.

Figure 88 on page 642 shows an example of the output produced by the SNMP query stack trace. This trace was produced by starting the SNMP query stack address space with start option -d 4, which is the maximum amount of trace records produced. In the figure, the column labeled “trc lvl” shows the lowest trace level required to see that particular trace entry. For example, lines five through nine have a “trc lvl” of four. This means that only the -d 4 trace option shows this type of trace entry. On the other hand, lines 10 through 17 have a “trc lvl” of two. This means that trace level two or higher produces this trace information.

Guideline: The column headed “line no.” numbers the trace records for reference in the discussion that follows the figure. Neither the “trc lvl” nor the “line no.” column appear in the actual trace output.

The following sequence of events occurred to create the trace output:

1. Started the SNMP query stack address space
   Trace output lines in the range 1–3
2. Started the SNMPIUCV subtask at the NetView host (attempted connection to the query stack when started)
   Trace output line 4
3. Issued an SNMP TRAPSON request (request 1001)
   Trace output lines in the range 5–27
4. Incoming SNMP Trap-PDU received from SNMP agent
   Trace output lines in the range 28–61
5. Issued an SNMP TRAPSOFF request (request 1002)
   Trace output lines in the range 62–82
6. Incoming SNMP Trap-PDU received from SNMP agent
   Trace output lines in the range 83–104
7. Issued an SNMP GET request (request 1003)
   Trace output lines in the range 105–148
8. Received the response to request 1003
   Trace output lines in the range 149–191
9. Issued an SNMP GETNEXT request (request 1004)

Trace output lines in the range 192–235

10. Received the response to request 1004
    Trace output lines in the range 236–278
11. Issued an SNMP SET request (request 1005)
    Trace output lines in the range 279–326
12. Received the response to request 1005
    Trace output lines in the range 327–369
13. Issued an SNMP MIBVNAME request (request 1006)
    Trace output lines in the range 370–397
14. Issued an SNMP PING request (request 1007)
    Trace output lines in the range 398–429
15. Issued an SNMP GET request for a variable name not defined in the *hlq*.MIBDESC.DATA data set (request 1008)
    Trace output lines in the range 430–462
16. Stopped the SNMPIUCV subtask of the NetView program
    Trace output line 463

```
trc  line
lvl  no.

 3    1  EZA6322I Using  'TCPCS.mibdesc.data'  as MIB description file
 0    2  EZA6275I SNMP Query Stack running and awaiting queries...
 2    3  EZA6276I There are 56 client connections possible
 0    4  EZA6290I Accepted new client connection
 4    5  EZA6292I Received following NVquery packet:
      6  EZA6305I dumping packet of 19 bytes:
      7  00 11 01 01 01 02 06 00 00 03 e9 00 00 00 00 00
      8  00 00 00
 2    9  EZA6359I major version:  1
     10  EZA6360I minor version:  1
     11  EZA6361I release:        1
     12  EZA6363I native set:     EBCDIC
     13  EZA6364I packet type:    TRAP REQUEST
     14  EZA6394I  filter id:         1001
     15  EZA6396I  network mask:      0.0.0.0
     16  EZA6397I  desired network:   0.0.0.0
 2   17  EZA6359I major version:  1
     18  EZA6360I minor version:  1
     19  EZA6361I release:        1
     20  EZA6363I native set:     EBCDIC
     21  EZA6364I packet type:    RESPONSE
     22  EZA6367I  sequence id:       1001
     23  EZA6388I  major error:       0
     24  EZA6389I  minor error:       0
     25  EZA6390I  error index:       0
     26  EZA6391I  error text len:    9
     27  EZA6392I  error text:        no error
 4   28  EZA6301I Received following SNMP_trap packet:
     29  EZA6305I dumping packet of 43 bytes:
     30           30 29 02 01 00 04 04 4d 56 53 4c a4 1e 06 0a 2b
     31           06 01 04 01 02 02 01 02 04 40 04 09 43 72 25 02
     32           01 04 02 01 00 43 02 25 80 30 00
 3   33  EZA6424I Decoded SNMP PDU :
     34  {
     35     version version-1,
     36     community '4d56534c'H,
     37     data {
     38        trap {
     39           enterprise 1.3.6.1.4.1.2.2.1.2.4,
     40           agent-addr {
     41              internet '09437225'H
     42           },
     43           generic-trap authenticationFailure,
     44           specific-trap 0,
     45           time-stamp 9600,
     46           variable-bindings {}
     47        }
     48     }
     49  }
 4   50  EZA6359I major version:  1
     51  EZA6360I minor version:  1
     52  EZA6361I release:        1
     53  EZA6363I native set:     EBCDIC
```

Figure 88. SNMP query engine traces

```
     54  EZA6364I packet type:        TRAP
     55  EZA6394I  filter id:         1001
     56  EZA6395I  agent address:     9.67.114.37
     57  EZA6399I  generic trap:      4     ( 0X4 )
     58  EZA6400I  specific trap:     0     ( 0X0 )
     59  EZA6401I  time stamp:        9600
     60  EZA6402I  enterprise len:    22
     61  EZA6403I  enterprise:        1.3.6.1.4.1.2.2.1.2.4
```

```
4   62  EZA6292I Received following NVquery packet:
    63  EZA6305I dumping packet of 15 bytes:
    64  00 0d 01 01 01 02 07 00 00 03 ea 00 00 03 e9
2   65  EZA6359I major version:  1
    66  EZA6360I minor version:  1
    67  EZA6361I release:        1
    68  EZA6363I native set:     EBCDIC
    69  EZA6364I packet type:    TRAP UN-REQUEST
    70  EZA6367I  sequence id:       1002
    71  EZA6394I  filter id:         1001
2   72  EZA6359I major version:  1
    73  EZA6360I minor version:  1
    74  EZA6361I release:        1
    75  EZA6363I native set:     EBCDIC
    76  EZA6364I packet type:    RESPONSE
    77  EZA6367I  sequence id:       1002
    78  EZA6388I  major error:       0
    79  EZA6389I  minor error:       0
    80  EZA6390I  error index:       0
    81  EZA6391I  error text len:    9
    82  EZA6392I  error text:        no error
4   83  EZA6301I Received following SNMP_trap packet:
    84  EZA6305I dumping packet of 43 bytes:
    85           30 29 02 01 00 04 04 4d 56 53 4c a4 1e 06 0a 2b
    86           06 01 04 01 02 02 01 02 04 40 04 09 43 72 25 02
    87           01 04 02 01 00 43 02 38 40 30 00
3   88  EZA6424I Decoded SNMP PDU :
    89  {
    90     version version-1,
    91     community '4d56534c'H,
    92     data {
    93        trap {
    94           enterprise 1.3.6.1.4.1.2.2.1.2.4,
    95           agent-addr {
    96              internet '09437225'H
    97           },
    98           generic-trap authenticationFailure,
    99           specific-trap 0,
   100           time-stamp 14400,
   101           variable-bindings {}
   102        }
   103     }
   104  }
4  105  EZA6292I Received following NVquery packet:
   106  EZA6305I dumping packet of 42 bytes:
   107  00 28 01 01 01 02 01 00 00 03 eb 00 05 d4 e5 e2
   108  d3 00 00 05 e2 d5 d4 d7 00 00 03 01 ff 01 00 0a
   109  e2 e8 e2 e4 d7 e3 c9 d4 c5 00
2  110  EZA6359I major version:  1
   111  EZA6360I minor version:  1
   112  EZA6361I release:        1
   113  EZA6363I native set:     EBCDIC
   114  EZA6364I packet type:    GET
   115  EZA6367I  sequence id:       1003
   116  EZA6368I  hostname len:      5
   117  EZA6370I  hostname:          MVSL
   118  EZA6371I  community len:     5
   119  EZA6373I  community:         SNMP
   120  EZA6374I  optional length:   3
   121  EZA6375I  max. retries:      1
   122  EZA6376I  initial timeout:   255
   123  EZA6377I  backoff exponent:  1
   124  EZA6380I   name length:      16
   125  EZA6381I   name:             1.3.6.1.2.1.1.3
3  126  EZA6424I Decoded SNMP PDU :
   127  {
   128     version version-1,
```

```
   129     community '534e4d50'H,
   130     data {
   131        get-request {
   132           request-id 1,
   133           error-status noError,
   134           error-index 0,
   135           variable-bindings {
   136              {
   137                 name 1.3.6.1.2.1.1.3,
   138                 value {
   139                    simple {
   140                       empty {}
   141                    }
   142                 }
   143              }
   144           }
   145        }
   146     }
   147  }
   148  EZA6308I sending SNMP request to 9.67.114.37
 4 149  EZA6295I Received following SNMP_response packet:
   150  EZA6305I dumping packet of 39 bytes:
   151           30 25 02 01 00 04 04 53 4e 4d 50 a2 1a 02 01 01
   152           02 01 00 02 01 00 30 0f 30 0d 06 07 2b 06 01 02
   153           01 01 03 43 02 48 a8
 3 154  EZA6424I Decoded SNMP PDU :
   155  {
   156     version version-1,
   157     community '534e4d50'H,
   158     data {
   159        get-response {
   160           request-id 1,
   161           error-status noError,
   162           error-index 0,
   163           variable-bindings {
   164              {
   165                 name 1.3.6.1.2.1.1.3,
   166                 value {
   167                    application-wide {
   168                       ticks 18600
   169                    }
   170                 }
   171              }
   172           }
   173        }
   174     }
   175  }
 2 176  EZA6359I major version:  1
   177  EZA6360I minor version:  1
   178  EZA6361I release:        1
   179  EZA6363I native set:     EBCDIC
   180  EZA6364I packet type:    RESPONSE
   181  EZA6367I  sequence id:       1003
   182  EZA6388I  major error:       0
   183  EZA6389I  minor error:       0
   184  EZA6390I  error index:       0
   185  EZA6391I  error text len:    9
   186  EZA6392I  error text:        no error
   187  EZA6380I   name length:      16
   188  EZA6381I   name:             1.3.6.1.2.1.1.3
   189  EZA6382I   value type:       time ticks
   190  EZA6384I   value length:     4
   191  EZA6387I   value:            18600
 4 192  EZA6292I Received following NVquery packet:
   193  EZA6305I dumping packet of 42 bytes:
   194           00 28 01 01 01 02 02 00 00 03 ec 00 05 d4 e5 e2
   195           d3 00 00 05 e2 d5 d4 d7 00 00 03 01 ff 01 00 0a
```

```
   196           c9 c6 c4 c5 e2 c3 d9 4b f0 00
2  197  EZA6359I major version:  1
   198  EZA6360I minor version:  1
   199  EZA6361I release:        1
   200  EZA6363I native set:     EBCDIC
   201  EZA6364I packet type:    GET-NEXT
   202  EZA6367I  sequence id:       1004
   203  EZA6368I  hostname len:      5
   204  EZA6370I  hostname:          MVSL
   205  EZA6371I  community len:     5
   206  EZA6373I  community:         SNMP
   207  EZA6374I  optional length:   3
   208  EZA6375I  max. retries:      1
   209  EZA6376I  initial timeout:   255
   210  EZA6377I  backoff exponent:  1
   211  EZA6380I   name length:      22
   212  EZA6381I   name:             1.3.6.1.2.1.2.2.1.2.0
3  213  EZA6424I Decoded SNMP PDU :
   214  {
   215     version version-1,
   216     community '534e4d50'H,
   217     data {
   218        get-next-request {
   219           request-id 2,
   220           error-status noError,
   221           error-index 0,
   222           variable-bindings {
   223              {
   224                 name 1.3.6.1.2.1.2.2.1.2.0,
   225                 value {
   226                    simple {
   227                       empty {}
   228                    }
   229                 }
   230              }
   231           }
   232        }
   233     }
   234  }
   235  EZA6308I sending SNMP request to 9.67.114.37
4  236  EZA6295I Received following SNMP_response packet:
   237  EZA6305I dumping packet of 47 bytes:
   238           30 2d 02 01 00 04 04 53 4e 4d 50 a2 22 02 01 02
   239           02 01 00 02 01 00 30 17 30 15 06 0a 2b 06 01 02
   240           01 02 02 01 02 01 04 07 49 42 4d 20 4c 43 53
3  241  EZA6424I Decoded SNMP PDU :
   242  {
   243     version version-1,
   244     community '534e4d50'H,
   245     data {
   246        get-response {
   247           request-id 2,
   248           error-status noError,
   249           error-index 0,
   250           variable-bindings {
   251              {
   252                 name 1.3.6.1.2.1.2.2.1.2.1,
   253                 value {
   254                    simple {
   255                       string '49424d204c4353'H
   256                    }
   257                 }
   258              }
   259           }
   260        }
   261     }
```

```
   262  }
 2 263  EZA6359I major version:  1
   264  EZA6360I minor version:  1
   265  EZA6361I release:        1
   266  EZA6363I native set:     EBCDIC
   267  EZA6364I packet type:    RESPONSE
   268  EZA6367I  sequence id:       1004
   269  EZA6388I  major error:       0
   270  EZA6389I  minor error:       0
   271  EZA6390I  error index:       0
   272  EZA6391I  error text len:    9
   273  EZA6392I  error text:        no error
   274  EZA6380I   name length:      22
   275  EZA6381I   name:             1.3.6.1.2.1.2.2.1.2.1
   276  EZA6382I   value type:       display string
   277  EZA6384I   value length:     7
   278  EZA6385I   value:            IBM LCS
 4 279  EZA6292I Received following NVquery packet:
   280  EZA6305I dumping packet of 57 bytes:
   281           00 37 01 01 01 02 03 00 00 03 ed 00 05 d4 e5 e2
   282           d3 00 00 05 e2 d5 d4 d7 00 00 03 01 ff 01 00 10
   283           c4 d7 c9 e2 c1 d4 d7 d3 c5 d5 e4 d4 c2 c5 d9 00
   284           00 00 06 f1 f2 f3 f4 f5 00
2  285  EZA6359I major version:  1
   286  EZA6360I minor version:  1
   287  EZA6361I release:        1
   288  EZA6363I native set:     EBCDIC
   289  EZA6364I packet type:    SET
   290  EZA6367I  sequence id:       1005
   291  EZA6368I  hostname len:      5
   292  EZA6370I  hostname:          MVSL
   293  EZA6371I  community len:     5
   294  EZA6373I  community:         SNMP
   295  EZA6374I  optional length:   3
   296  EZA6375I  max. retries:      1
   297  EZA6376I  initial timeout:   255
   298  EZA6377I  backoff exponent:  1
   299  EZA6380I   name length:      22
  300  EZA6381I   name:             1.3.6.1.4.1.2.2.1.4.1
   301  EZA6382I   value type:       number
   302  EZA6384I   value length:     4
   303  EZA6386I   value:            12345
 3 304  EZA6424I Decoded SNMP PDU :
   305  {
   306     version version-1,
   307     community '534e4d50'H,
   308     data {
   309        set-request {
   310           request-id 3,
   311           error-status noError,
   312           error-index 0,
   313           variable-bindings {
   314              {
   315                 name 1.3.6.1.4.1.2.2.1.4.1,
   316                 value {
   317                    simple {
   318                       number 12345
   319                    }
   320                 }
   321              }
   322           }
   323        }
   324     }
   325  }
   326  EZA6308I sending SNMP request to 9.67.114.37
 4 327  EZA6295I Received following SNMP_response packet:
   328  EZA6305I dumping packet of 42 bytes:
```

```
   329           30 28 02 01 00 04 04 53 4e 4d 50 a2 1d 02 01 03
   330           02 01 00 02 01 00 30 12 30 10 06 0a 2b 06 01 04
   331           01 02 02 01 04 01 02 02 30 39
3  332  EZA6424I Decoded SNMP PDU :
   333  {
   334     version version-1,
   335     community '534e4d50'H,
   336     data {
   337        get-response {
   338           request-id 3,
   339           error-status noError,
   340           error-index 0,
   341           variable-bindings {
   342              {
   343                 name 1.3.6.1.4.1.2.2.1.4.1,
   344                 value {
   345                    simple {
   346                       number 12345
   347                    }
   348                 }
   349              }
   350           }
   351        }
   352     }
   353  }
2  354  EZA6359I major version:  1
   355  EZA6360I minor version:  1
   356  EZA6361I release:        1
   357  EZA6363I native set:     EBCDIC
   358  EZA6364I packet type:    RESPONSE
   359  EZA6367I  sequence id:       1005
   360  EZA6388I  major error:       0
   361  EZA6389I  minor error:       0
   362  EZA6390I  error index:       0
   363  EZA6391I  error text len:    9
   364  EZA6392I  error text:        no error
   365  EZA6380I   name length:      22
   366  EZA6381I   name:             1.3.6.1.4.1.2.2.1.4.1
   367  EZA6382I   value type:       number
   368  EZA6384I   value length:     4
   369  EZA6386I   value:            12345
4  370  EZA6292I Received following NVquery packet:
   371  EZA6305I dumping packet of 29 bytes:
   372           00 1b 01 01 01 02 08 00 00 03 ee 00 10 f1 4b f3
   373           4b f6 4b f1 4b f2 4b f1 4b f1 4b f1 00
2  374  EZA6359I major version:  1
   375  EZA6360I minor version:  1
   376  EZA6361I release:        1
   377  EZA6363I native set:     EBCDIC
   378  EZA6364I packet type:    VAR_NAME
   379  EZA6367I  sequence id:       1006
   380  EZA6405I  object id len:     16
   381  EZA6406I  object id:         1.3.6.1.2.1.1.1
2  382  EZA6359I major version:  1
   383  EZA6360I minor version:  1
   384  EZA6361I release:        1
   385  EZA6363I native set:     EBCDIC
   386  EZA6364I packet type:    RESPONSE
   387  EZA6367I  sequence id:       1006
   388  EZA6388I  major error:       0
   389  EZA6389I  minor error:       0
   390  EZA6390I  error index:       0
   391  EZA6391I  error text len:    9
   392  EZA6392I  error text:        no error
   393  EZA6380I   name length:      16
   394  EZA6381I   name:             1.3.6.1.2.1.1.1
   395  EZA6382I   value type:       display string
```

```
   396  EZA6384I   value length:     9
   397  EZA6385I   value:            sysDescr
 4 398  EZA6292I Received following NVquery packet:
   399  EZA6305I dumping packet of 23 bytes:
   400           00 15 01 01 01 02 0a 00 00 03 ef 00 05 d4 e5 e2
   401           d3 00 00 03 01 ff 01
 2 402  EZA6359I major version:  1
   403  EZA6360I minor version:  1
   404  EZA6361I release:        1
   405  EZA6363I native set:     EBCDIC
   406  EZA6364I packet type:    PING REQUEST
   407  EZA6367I  sequence id:       1007
   408  EZA6368I  hostname len:      5
   409  EZA6370I  hostname:          MVSL
   410  EZA6374I  optional length:   3
   411  EZA6375I  max. retries:      1
   412  EZA6376I  initial timeout:   255
   413  EZA6377I  backoff exponent:  1
 2 414  EZA6359I major version:  1
   415  EZA6360I minor version:  1
   416  EZA6361I release:        1
   417  EZA6363I native set:     EBCDIC
   418  EZA6364I packet type:    RESPONSE
   419  EZA6367I  sequence id:       1007
   420  EZA6388I  major error:       0
   421  EZA6389I  minor error:       0
   422  EZA6390I  error index:       0
   423  EZA6391I  error text len:    9
   424  EZA6392I  error text:        no error
   425  EZA6380I   name length:      34
   426  EZA6381I   name:             1.3.6.1.4.1.2.2.1.3.2.9.67.114.37
   427  EZA6382I   value type:       number
   428  EZA6384I   value length:     4
   429  EZA6386I   value:            76
 4 430  EZA6292I Received following NVquery packet:
   431  EZA6305I dumping packet of 39 bytes:
   432           00 25 01 01 01 02 01 00 00 03 f0 00 05 d4 e5 e2
   433           d3 00 00 05 e2 d5 d4 d7 00 00 03 01 ff 01 00 07
   434           c2 c1 c4 e5 c1 d9 00
 1 435  EZA6356E error code 7: unknown MIB variable
 2 436  EZA6359I major version:  1
   437  EZA6360I minor version:  1
   438  EZA6361I release:        1
   439  EZA6363I native set:     EBCDIC
   440  EZA6364I packet type:    GET
   441  EZA6367I  sequence id:       1008
   442  EZA6368I  hostname len:      5
   443  EZA6370I  hostname:          MVSL
   444  EZA6371I  community len:     5
   445  EZA6373I  community:         SNMP
   446  EZA6374I  optional length:   3
   447  EZA6375I  max. retries:      1
   448  EZA6376I  initial timeout:   255
   449  EZA6377I  backoff exponent:  1
   450  EZA6380I   name length:      2
   451  EZA6381I   name:             ?
2  452  EZA6359I major version:  1
   453  EZA6360I minor version:  1
   454  EZA6361I release:        1
   455  EZA6363I native set:     EBCDIC
   456  EZA6364I packet type:    RESPONSE
   457  EZA6367I  sequence id:       1008
   458  EZA6388I  major error:       2
   459  EZA6389I  minor error:       7
```

```
   460  EZA6390I  error index:       0
   461  EZA6391I  error text len:    17
   462  EZA6392I  error text:        unknown variable
0  463  EZA6293I Terminated client connection
```

The following is an explanation of the traces in Figure 88 on page 642.

- Line 1 is an information message listing the actual name of the data set being used as the *hlq*.MIBDESC.DATA data set.
- Line 2 is an informational message indicating that the SNMP query stack has been successfully started.
- Line 3 is an informational message indicating the number of client connections the query stack allows. (A client connection is a connection from a program using the query stack protocol to communicate with the SNMP query stack to initiate SNMP requests. For example, the SNMPIUCV subtask of the NetView program is a client connection).
- Line 4 is an information message indicating that the SNMPIUCV subtask of the NetView program has successfully contacted the query stack.
- Lines 5–8 are the encoded packet received from the client (the SNMPIUCV subtask) by the query stack. This particular packet is the TRAPSON request.
- Lines 9–16 are the decoded SNMPIUCV request. The decoded packet indicates that this request is number 1001 (line 14), and was a TRAPSON request (line 13) for network mask 0.0.0.0 (line 15) with the desired network 0.0.0.0 (line 16).
- Lines 17–27 are the response sent back to SNMPIUCV from the query stack. The response (line 21) is to request number 1001 (line 22) and indicates that the TRAPSON request was successful (lines 23–27).
- Line 28 indicates that an SNMP Trap-PDU was received. Lines 29–32 are the actual BER encoded SNMP packet as it was received by the query stack.
- Lines 33–49 are the decoded version of the trap packet reported by lines 28–32.
- Lines 50–61 are the trap information being passed from the query stack up to the SNMPIUCV subtask to be displayed to the NetView operator. This trap is being forwarded to the NetView program because the IP address of the agent sending the trap (line 56), when ANDed with the network mask (line 15) matches the desired network (line 16) of filter number 1001 (line 55) that was set by the TRAPSON request 1001 (line 14) received previously (lines 9–16).
- Lines 62–64 show an incoming query stack packet sent from SNMPIUCV to the query stack.
- Lines 65–71 are the decoded packet received in lines 62–64. This packet is the TRAPSOFF request (line 69). It requests that trap filter 1001 (line 71) be turned off.
- Lines 72–82 are the response from the query stack to the SNMPIUCV subtask. The response indicates that the TRAPSOFF request was completed successfully (lines 78–82).
- Lines 83–87 indicate that another SNMP Trap-PDU was received from an agent.
- Lines 88–104 are the decoded Trap-PDU. Note that following this decoded PDU, there is no indication of the trap being forwarded to SNMPIUCV. This is because the trap filter has been turned off, so the query stack receives the trap but does not forward the information to SNMPIUCV.
- Lines 105–109 indicate another request from SNMPIUCV being received by the query stack.
- Lines 110–125 are the decoded query stack request. The request from SNMPIUCV was to issue a GetRequest-PDU (line 114) to host MVSL (line 117), using community name SNMP (line 119) and requesting variable 1.3.6.1.2.1.1.3

(line 125). Lines 121–123 are the retry information that SNMPIUCV has gotten from the SNMPARMS member of the DSIPARMS data set.

- Lines 126–147 are the decoded SNMP GetRequest-PDU that the query stack has built as a result of the SNMPIUCV request received in lines 110–125. This PDU has been assigned request number 1 (line 132). This number is used to correlate the response when it is received.
- Line 148 indicates that the encoded SNMP GetRequest-PDU has been sent to the SNMP agent at the specified IP address. This should be the IP address of the host specified in line 117.
- Line 149 indicates that an SNMP GetResponse-PDU was received. Lines 150–153 are the encoded GetResponse-PDU.
- Lines 154–175 are the decoded GetResponse-PDU. This was a GetResponse (line 159) in response to request number 1 (line 160, matches up to the request number in the request, line 132). The request was completed with no errors (lines 161–162). The requested variable 1.3.6.1.2.1.1.3 (line 165) has a value of 18600 timeticks (line 168).
- Lines 176–191 are the query stack response to SNMPIUCV request number 1003 (lines 115 and 181). The response contains the information received from the agent in the GetResponse-PDU in lines 154–175.
- Lines 192–196 are the next query stack protocol requests received from SNMPIUCV by the query stack.
- Lines 197–212 are the decoded version of the query stack request. This is a GetNext request (line 201) to host MVSL (line 204) for variable 1.3.6.1.2.1.2.2.1.2.0 (line 212). The request number associated with this request is 1004 (line 202).
- Lines 213–234 are the decoded SNMP GetRequest-PDU built as a result of the query stack request received in lines 197–212. This GetRequest-PDU is request number 2 (line 219).
- Line 235 indicates that the encoded GetRequest-PDU has been sent to the requested host.
- Lines 236–240 indicate that a GetResponse-PDU has been received.
- Lines 241–262 are the decoded GetResponse-PDU. This is the response to request number 2 (line 247) for variable 1.3.6.1.2.1.2.2.1.2.1 (line 252). The value of the variable is a display string with the ASCII value of X'49424D204C4353' (line 255). The GetNext request completed successfully (lines 248–249).
- Lines 263–278 are the query stack response to SNMPIUCV request 1004 (line 268). The response contains the information from the GetResponse-PDU (lines 241–262). Note that the variable value in line 255 has been converted to the proper display format in line 278.
- Lines 279–284 are the next query stack protocol request from SNMPIUCV to the query stack.
- Lines 285–303 are the decoded query stack request. It is a SET request (line 289) to host MVSL to set variable 1.3.6.1.4.1.2.2.1.4.1 (line 300) to 12345 (line 303). This is request number 1005 (line 290).
- Lines 304–325 are the decoded SNMP SetRequest-PDU built as a result of the request received in lines 285–303. This is request number 3 (line 310).
- Line 326 indicates that the SetRequest-PDU has been sent to the specified host.
- Lines 327–331 indicate that a GetResponse-PDU has been received.
- Lines 332–353 are the decoded GetResponse-PDU. This PDU is the response to the SetRequest-PDU number 3 (line 338). It was completed successfully (lines 339–340) and variable 1.3.6.1.4.1.2.2.1.4.1 (line 343) was set to 12345 (line 346).

- Lines 354–369 are the query stack response to request 1005 (line 359) containing the information received in the GetResponse-PDU received in lines 332–353.
- Lines 370–373 are the next query stack request packet from SNMPIUCV.
- Lines 374–381 are the decoded query stack request. This is a MIBVNAME request (line 378) requesting the name of variable 1.3.6.1.2.1.1.1 (line 381). The request number is 1006 (line 379).
- Lines 382–397 are the query stack response (line 386) to request 1006 (line 387). The request completed successfully (lines 388–392) and the name of variable 1.3.6.1.2.1.1.1 (line 394) is sysDescr (line 397).
- Lines 398–401 are the next query stack request packet from SNMPIUCV.
- Lines 402–413 are the decoded query stack request packet. This is a PING request (line 406) to ping host MVSL (line 409). The request number is 1007 (line 407).
- Lines 414–429 are the query stack response (line 418) to request 1007 (line 419). The PING completed successfully (lines 420–424) and the round-trip response time was 76 milliseconds (line 429). Note that no SNMP PDUs were generated as a result of the PING request. The SNMP query stack uses a raw socket to send a PING to the requested host and SNMP protocols are not involved.
- Lines 430–434 are the next query stack request packet received from SNMPIUCV.
- Line 435 indicates that an error occurred while the query stack was decoding the request packet. The MIB variable name in the request was unknown to the query stack.
- Lines 435–451 are the decoded query stack request. This is a GET request (line 440). The variable name is unknown (line 451). This is request 1008 (line 441).
- Lines 452–462 are the query stack response (line 456) to SNMPIUCV request 1008 (line 457). The request was unsuccessful. The query stack returns major error code 2 (line 458), minor error code 7 (line 459), unknown variable (line 462). Note that no SNMP PDUs were generated since the query stack could not resolve the variable name.
- Line 463 indicates that the client connection (SNMPIUCV) has been terminated. This is the result of the STOP TASK=SNMPIUCV command.

NetView SNMP query engine IUCV communication trace

Figure 89 on page 652 shows an example of the output produced by the IUCV communication trace. This trace was produced by starting the SNMP query stack address space with start option -it.

```
descarray is at 3985ab8, size is 4 bytes
descarray has 50 entries, entry size is 928
iucvdesc is at 32508
Rc=0 on IUCV_CLR to TCPCS   , fd=-254, path=0, iprcode=0, ipmsgid=0, iucvname=00032508
     ciucv_data area (ipbfadr2) is at 00000000
Rc=0 on IUCV_SET to TCPCS   , fd=-254, path=0, iprcode=0, ipmsgid=0, iucvname=00032508
     ciucv_data area (ipbfadr2) is at 00005480
Rc=0 on IUCV_CONNECT to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=A0000,
iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00005494
     IUCV interrupt from TCPIP, fd=-254, path=1 type=2 (Connection Complete)
sock_request_inet entry parms:
   f=0 d=-254 rl=00000000 rd=0005ddfc rdl=20 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=C, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 000054bc
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=0 d=-254 rl=00000000 rd=0005ddfc rdl=20 pdh=0 pdl=0
   rc=0 err=49 rpl=00000000 rpb=00000000 rpbl=0
sock_request_inet entry parms:
   f=25 d=3 rl=00000000 rd=0005db2c rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=3, path=1, iprcode=0, ipmsgid=D, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=3 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=3 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=3 buf (ipbfadr2) is at 000054e4
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=25 d=3 rl=00000000 rd=0005db2c rdl=16 pdh=0 pdl=0
   rc=3 err=0 rpl=00000000 rpb=00000000 rpbl=0
sock_request_inet entry parms:
   f=2 d=3 rl=00000000 rd=0001d0d8 rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=3, path=1, iprcode=0, ipmsgid=E, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=3 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=3 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=3 buf (ipbfadr2) is at 0000550c
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=2 d=3 rl=00000000 rd=0001d0d8 rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
sock_request_inet entry parms:
   f=25 d=4 rl=00000000 rd=0005db44 rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=4, path=1, iprcode=0, ipmsgid=F, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 00005534
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
```

Figure 89. SNMP IUCV communication traces

```
sock_request_inet return parms:
   f=25 d=4 rl=00000000 rd=0005db44 rdl=16 pdh=0 pdl=0
   rc=4 err=0 rpl=00000000 rpb=00000000 rpbl=0
sock_request_inet entry parms:
   f=2 d=4 rl=00000000 rd=0005da9c rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
```

```
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=4, path=1, iprcode=0, ipmsgid=10, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 0000555c
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=2 d=4 rl=00000000 rd=0005da9c rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
sock_request_inet entry parms:
   f=25 d=5 rl=00000000 rd=0005db64 rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=5, path=1, iprcode=0, ipmsgid=11, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 00005584
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=25 d=5 rl=00000000 rd=0005db64 rdl=16 pdh=0 pdl=0
   rc=5 err=0 rpl=00000000 rpb=00000000 rpbl=0
sock_request_inet entry parms:
   f=2 d=5 rl=00000000 rd=0005daa8 rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=5, path=1, iprcode=0, ipmsgid=12, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 000055ac
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=2 d=5 rl=00000000 rd=0005daa8 rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
sock_request_inet entry parms:
   f=13 d=5 rl=00000000 rd=00000000 rdl=0 pdh=0 pdl=5
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=5, path=1, iprcode=0, ipmsgid=13, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=5 buf (ipbfadr2) is at 000055d4
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=13 d=5 rl=00000000 rd=00000000 rdl=0 pdh=0 pdl=5
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_SET to , fd=6, path=0, iprcode=0, ipmsgid=0, iucvname=SNMPQE
     ciucv_data area (ipbfadr2) is at 00005480
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
sock_request_inet entry parms:
   f=25 d=7 rl=00000000 rd=0005db2c rdl=16 pdh=0 pdl=0
   rc=0 err=0 rpl=00000000 rpb=00000000 rpbl=0
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=7, path=1, iprcode=0, ipmsgid=14, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=7 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=7 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=7 buf (ipbfadr2) is at 000055fc
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=25 d=7 rl=00000000 rd=0005db2c rdl=16 pdh=0 pdl=0
   rc=7 err=0 rpl=00000000 rpb=00000000 rpbl=0
 SQEI001 -- SNMP Query Engine running and awaiting queries...
fd=3 in callers rmask
fd=4 in callers rmask
fd=5 in callers rmask
fd=6 in callers rmask
fd=7 in callers rmask
```

```
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
fd=3 inetselect now TRUE
fd=6 iucvselect now TRUE
in inetselect
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=15, iucvname=00032508
wait ecblist=5dc5c, ecbcount=2
iucvposted=1073741824, waitposted=0, callposted=0
in iucvposted
Rc=1 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00005624
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
in gotmsgcomp
fd=3 inetselect now TRUE
fd=6 iucvselect now TRUE
nfds=0, return=1
sock_request_inet entry parms:
   f=16 d=4 rl=00000000 rd=00000000 rdl=0 pdh=0 pdl=0
   rc=0 err=0 rpl=0005ed88 rpb=00000000 rpbl=4120
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=4, path=1, iprcode=0, ipmsgid=16, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 00000000
Rc=1 on IUCV_NEXTBUFF, fd=4 buf (ipbfadr2) is at 0000564c
     IUCV interrupt from TCPIP, fd=-254, path=1 type=7 (Incoming Reply)
sock_request_inet return parms:
   f=16 d=4 rl=00000000 rd=00000000 rdl=0 pdh=0 pdl=0
   rc=0 err=0 rpl=0005ed88 rpb=00000000 rpbl=68
fd=3 in callers rmask
fd=4 in callers rmask
fd=5 in callers rmask
fd=6 in callers rmask
fd=7 in callers rmask
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
fd=3 inetselect now TRUE
fd=6 iucvselect now TRUE
in inetselect
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=17, iucvname=00032508
wait ecblist=5dc5c, ecbcount=2
iucvposted=1073741824, waitposted=0, callposted=0
in iucvposted
Rc=1 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00005674
     IUCV interrupt, fd=6, path=2 type=1 (Pending Connection)
iucvcomp is now TRUE
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
in iucvcom && iucvselect
fd=3 inetselect now TRUE
fd=6 iucvselect now TRUE
Rc=0 on IUCV_PURGE to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=17, iucvname=00032508
Rc=0 on IUCV_NEXTBUFF, fd=6 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_ACCEPT to CNMR3X  , fd=8, path=2, iprcode=0, ipmsgid=10000, iucvname=SNMPQE
 SQEI002 -- Accepted new client connection
fd=3 in callers rmask
fd=4 in callers rmask
fd=5 in callers rmask
fd=6 in callers rmask
fd=7 in callers rmask
fd=8 in callers rmask
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
fd=3 inetselect now TRUE
fd=6 iucvselect now TRUE
in inetselect
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=18, iucvname=00032508
wait ecblist=5dc5c, ecbcount=2
iucvposted=1073741824, waitposted=0, callposted=0
```

```
in iucvposted
Rc=1 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 0000569c
     IUCV interrupt, fd=8, path=2 type=3 (Connection Severed)
iucvcomp is now TRUE
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
in iucvcom && iucvselect
fd=3 inetselect now TRUE
fd=6 iucvselect now TRUE
Rc=0 on IUCV_PURGE to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=18, iucvname=00032508
fd=8 in callers rmask
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
fd=8 iucvselect now TRUE
in iucvselect, iucvnfds=1
Rc=0 on IUCV_SEVER to CNMR3X  , fd=8, path=2, iprcode=0, ipmsgid=0, iucvname=SNMPQE
 SQEI003 -- Terminated client connection
fd=3 in callers rmask
fd=4 in callers rmask
fd=5 in callers rmask
fd=6 in callers rmask
fd=7 in callers rmask
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
fd=3 inetselect now TRUE
fd=6 iucvselect now TRUE
in inetselect
Rc=0 on IUCV_NEXTBUFF, fd=-254 buf (ipbfadr2) is at 00000000
Rc=0 on IUCV_SEND to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=19, iucvname=00032508
wait ecblist=5dc5c, ecbcount=2
iucvposted=0, waitposted=0, callposted=1073741824
callers ECB posted
Rc=0 on IUCV_PURGE to TCPCS   , fd=-254, path=1, iprcode=0, ipmsgid=19, iucvname=00032508
Rc=0 on IUCV_CLR to , fd=6, path=0, iprcode=0, ipmsgid=0, iucvname=SNMPQE
     ciucv_data area (ipbfadr2) is at 00005480
Rc=0 on IUCV_CLR to TCPCS   , fd=-254, path=0, iprcode=0, ipmsgid=0, iucvname=00032508
     ciucv_data area (ipbfadr2) is at 00000000
```

The following sequence of events occurred to create the trace output:

1. Started the SNMP query stack
2. Connected to the query stack from the SNMPIUCV subtask
3. Disconnected the SNMPIUCV subtask from the query stack

TRAPFWD Trace

The trap forwarder daemon uses syslog functions to write out debug information and traces. Diagnostic data is written using "trapfwd" as identifier.

Figure 90 on page 656 illustrates a TRAPFWD trace.

```
Oct 15 14:06:06  trapfwd[16777250]: EZZ8420I The Trap Forwarder daemon is running as USER17
Oct 15 14:06:06  trapfwd[16777250]: Establishing affinity with the TCPIP stack
Oct 15 14:06:06  trapfwd[16777250]: Issuing setibmopt for TCPCS
Oct 15 14:06:06  trapfwd[16777250]: Checking if TCP/IP stack is enabled
Oct 15 14:06:06  trapfwd[16777250]: Reading the configuration file : /etc/trapfwd.conf
Oct 15 14:06:06  trapfwd[16777250]: Line 1 : 9.67.113.79 2162
Oct 15 14:06:06  trapfwd[16777250]: Added entry with host: 9.67.113.79 port: 2162
Oct 15 14:06:06  trapfwd[16777250]: Line 2 : 9.67.113.79 1062
Oct 15 14:06:06  trapfwd[16777250]: Added entry with host: 9.67.113.79 port: 1062
Oct 15 14:06:06  trapfwd[16777250]: Line 3 : 9.67.113.79 169
Oct 15 14:06:06  trapfwd[16777250]: Added entry with host: 9.67.113.79 port: 169
Oct 15 14:06:06  trapfwd[16777250]: Line 4 : 9.67.113.79 179
Oct 15 14:06:06  trapfwd[16777250]: Added entry with host: 9.67.113.79 port: 179
Oct 15 14:06:06  trapfwd[16777250]: Creating sockets...
Oct 15 14:06:07  trapfwd[16777250]: EZZ8409I TRAPFWD: INITIALIZATION COMPLETE
Oct 15 14:06:07  trapfwd[16777250]: Ready to receive and forward traps....
```

Figure 90. TRAPFWD trace

Chapter 25. Diagnosing Policy Agent problems

The z/OS UNIX Policy Agent (PAGENT) provides administrative control for policies. This topic provides information and guidance to diagnose Policy Agent problems, and it contains the following subtopics:

- "Overview"
- "QoS policy" on page 658
- "QoS policy scope" on page 658
- "Gathering diagnostic information" on page 659
- "Diagnosing Policy Agent problems" on page 661

Overview

The Policy Agent can act in any of several roles, depending on configuration options:

- The Policy Agent can act as the Policy Decision Point (PDP) on a single system, installing policies in one or more z/OS Communications Server stacks.
- The Policy Agent can act as a centralized *policy server*, providing PDP services for one or more remote policy clients.
- The Policy Agent can act as a *policy client*, retrieving remote policies from the policy server. Each stack in a Common INET (CINET) environment that is configured to the Policy Agent acts as a separate policy client.
- A single Policy Agent can act as a policy client or a policy server, but not both.

Policy Agent reads policies defined in local or remote configuration files, or reads by way of the Lightweight Directory Access Protocol (LDAP) from an LDAP server. These policies are then installed in one or more TCP/IP stacks. Policy Agent can be configured to install identical policies to multiple (or all) stacks, or can install different sets of policies to each stack individually. Policy Agent can also monitor its configuration files and the LDAP server periodically for changed policies, and install new or changed policies as changes occur. The basic types of policies are:

- Quality of Service (QoS)
- Intrusion Detection Services (IDS)

  See Chapter 27, "Diagnosing intrusion detection problems," on page 699 for more information about diagnosing IDS policies.
- IPSec

  See Chapter 30, "Diagnosing IP security and defensive filter problems," on page 733 for more information about diagnosing IPSec policies.
- Application Transparent Transport Layer Security (AT-TLS)

  See Chapter 28, "Diagnosing Application Transparent Transport Layer Security (AT-TLS)," on page 707 for more information about diagnosing AT-TLS policies.
- Policy-based routing (Routing)

  See "Steps for diagnosing problems with IP routing to a destination when using policy-based routing" on page 31 for more information about diagnosing routing policies.

See z/OS Communications Server: IP Configuration Guide for more information about configuring and starting Policy Agent, as well as defining policies.

© Copyright IBM Corp. 2000, 2015

QoS policy

You need to become familiar with the following terms to understand QoS policies:

Quality of Service (QoS)
The overall service that a user or application receives from a network, in terms of throughput, delay, and such

Service Differentiation
The ability of a network to provide different QoS levels to different users or applications based on their needs.

Service Level Agreement (SLA)
A contract, in business terms, provided by a network service provider that details the QoS that users or applications are expected to receive.

Service Policy
Administrative controls for a network, which are needed to achieve the QoS promised by a given SLA.

Integrated Services
A type of service that provides end-to-end QoS to an application, using the methodology of resource reservation along the data path from a receiver to a sender.

Differentiated Services
A type of service that provides QoS to broad classes of traffic or users, for example, all FTP traffic to a given subnet.

Resource ReSerVation Protocol (RSVP)
A protocol that provides for resource reservation in support of Integrated Services.

QoS policy scope

QoS policies can be defined with different scopes. The following scopes are supported:

DataTraffic
The policy applies to generic data traffic. This type of policy is in support of Differentiated Services.

RSVP The policy applies to RSVP data traffic. This type of policy is in support of Integrated Services.

The TCP/IP stack maps TCP, UDP, and RAW traffic to QoS policies based on the selection criteria defined in the policy. Search criteria can include, but are not limited to, items such as source and destination IP addresses and ports, protocol, and interfaces. The mapping of DataTraffic scoped policies occurs at connect time for TCP traffic, and for each packet for UDP and RAW traffic. However, for UDP and RAW, the mappings are cached such that subsequent packets sent to the same destination use the cached mapping. RSVP scoped policies are only mapped when the RSVP Agent adds a reservation requested by an RSVP application. The mapping is removed when the reservation is removed. See Chapter 26, “Diagnosing RSVP agent problems,” on page 685 for more information about the operation of RSVP.

You can see the effect of defined QoS policies in the following ways:

- Use the Network SLAPM2 Subagent to display service policy and mapped application information, as well as to manage and display Network SLAPM2 performance monitoring.
- Use the z/OS UNIX **pasearch**, z/OS UNIX **netstat**, and TSO NETSTAT commands as follows:
  - The **pasearch** command shows defined policies.
  - The NETSTAT SLAP or **netstat -j** command shows performance metrics for active QoS policy rules.
  - The NETSTAT ALL or **netstat -A** command has additional information for each active connection that shows the QoS policy rule name if the connection maps to a QoS policy.

See the z/OS Communications Server: IP System Administrator's Commands for more information about the Netstat command, the **pasearch** command, and the Network SLAPM2 Subagent.

Import services

The IBM Configuration Assistant for z/OS Communications Server can request existing information from the Policy Agent. When the Policy Agent is providing this import service, the IBM Configuration Assistant is acting as an import requestor. The import requestor can make the following kinds of requests:

- Request that existing policy configuration files be imported for further changes and additions. These policy files are called import configuration files; the resulting policies are called import policies.
- Request that TCP/IP profile information be imported. This profile information can be used by the import requestor as configuration data.

See the Policy-based networking information in z/OS Communications Server: IP Configuration Guide for more details.

Gathering diagnostic information

Policy Agent writes logging information to a log file. The level of logged information is controlled by the LogLevel configuration statement and the -d startup option. This information (loglevel and debug level) can also be changed after startup using the MODIFY command as shown in the following example:

```
MODIFY procname,LOGLEVEL,LEVEL=127
MODIFY procname,DEBUG,LEVEL=2
```

Error, console, warning, and event LogLevel messages are written by default. To gather more diagnostic information, you can specify a LogLevel value greater than the default or specify debug level 1. This debug level has the side effect of setting the maximum LogLevel value as well.

If you are using both a policy server and one or more policy clients, be sure to gather the log files from all affected Policy Agent applications.

Use the debug levels as follows:

Debug level 1

Use debug level 1 for most debugging, except Sysplex Distributor performance monitor. This debug value gives extra debugging messages and uses the maximum LogLevel for logging.

Debug level 2
Use debug level 2 to verify Policy Agent processing of LDAP objects, or if a problem is suspected in how LDAP objects are defined.

Debug level 4
Use debug level 4 for summary information concerning Sysplex Distributor performance monitor QoS fraction calculations.

Debug level 8
Use debug level 8 for detailed information concerning Sysplex Distributor performance monitor QoS fraction calculations, and additional Sysplex Distributor debugging.

Debug level 16
Use debug level 16 to assist with memory allocation and leak problems. This debug value causes memory allocation and free requests to be logged inline. This can be used in conjunction with the -m startup option and the MODIFY MEMTRC command to debug memory problems.

Debug level 32
Use debug level 32 for detailed information about all policies as they are installed in the TCP/IP stack.

Debug level 64
Use debug level 64 for detailed locking information within Policy Agent.

Debug level 128
Use debug level 128 for details about remote PAPI connections on the policy server, and about connections to the policy server on the policy client.

Debug level 256
Use debug level 256 for details about requests to discover TCP/IP profile information from import requestors.

Use the trace option -t to turn on LDAP client library debugging. Use the trace levels as follows:

Trace level 0
Use trace level 0 for no LDAP client library debugging. This is the default.

Trace level 1
Use trace level 1 to turn on LDAP client library debugging. Note that the destination of LDAP client debug messages is *stderr*. This is controlled by the LDAP client library, not Policy Agent. Using trace level 1 turns on the following LDAP DEBUG options:

- LDAP_DEBUG_TRACE
- LDAP_DEBUG_PACKETS
- LDAP_DEBUG_ARGS
- LDAP_DEBUG_CONNS
- LDAP_DEBUG_BER
- LDAP_DEBUG_FILTER
- LDAP_DEBUG_MESSAGE
- LDAP_DEBUG_STATS
- LDAP_DEBUG_THREAD
- LDAP_DEBUG_PARSE
- LDAP_DEBUG_PERFORMANCE

- LDAP_DEBUG_REFERRAL
- LDAP_DEBUG_ERROR

Trace option disabled
If you start Policy Agent with the trace option disabled, the stderr output destination is closed.

Restriction: You cannot turn on the trace option later with the MODIFY command.

See z/OS Communications Server: IP Configuration Reference for details on how to use the LogLevel, debug level, and trace level.

Log output can be directed either to a set of log files or to the syslog daemon (syslogd). This can be accomplished with the `-l` startup option or the PAGENT_LOG_FILE environment variable. If output is directed to log files, the number and size of the files can be controlled using the PAGENT_LOG_FILE_CONTROL environment variable. This environment variable can be used to extend the size of the log information collected if necessary. For example, if a large LDAP configuration is used with debug level 2, the default log file size and number might not be sufficient to capture all of the information needed. In this case, use the environment variable to increase the number or size, or the number and size, of the log files. See z/OS Communications Server: IP Configuration Guide for more details on using LogLevel, the -d startup option, and the environment variables, as well as the location of the log file.

The following additional information might be useful in diagnosing Policy Agent problems:

- Output from the **pasearch** command
- Output from the NETSTAT IDS or **netstat -k** commands
- Output from the NETSTAT SLAP or **netstat -j** commands
- Output from the NETSTAT ALL or **netstat -A** commands for active connections mapped to policies
- Output from the **ipsec** command for IPSec policies
- Output from the NETSTAT TTLS or **netstat -x** command for AT-TLS policies
- SNMP output from walks of the Network SLAPM2 subagent MIB tables
- TCP/IP CTRACE output, using the POLICY, INTERNET and IOCTL CTRACE options
- RSVP Agent log output if RSVP scoped policies are defined

Diagnosing Policy Agent problems

Policy Agent problems generally fall into one of the following categories:

- "Initialization problems" on page 662
- "Policy definition problems" on page 663
- "Policy client connection problems" on page 668
- "Policy client retrieval problems" on page 671
- "Import requestor connection problems" on page 673
- "Import requestor retrieval problems" on page 678
- "LDAP object retrieval problems" on page 678
- "LDAP object storage problems" on page 680
- "Policy Agent and Sysplex distribution problems" on page 681

• "Memory allocation/leakage problems" on page 682

Initialization problems

If Policy Agent does not complete initialization, or fails to install any policies in one or more stacks, run it with the -d 1 startup option, and check the log file for error conditions. If Policy Agent fails to initialize, message EZZ8434I is issued to the console. Check the log file for the specific error encountered.

Table 48 lists some common Policy Agent initialization problems.

Table 48. Common Policy Agent initialization problems

Problem	Cause/action	Symptom
Policy Agent started from a user ID without superuser authority	Policy Agent must be started from a superuser	EZZ8434I message, along with messages in the log file indicating that superuser authority is required and showing an exit code value of 27.
Policy Agent not authorized to security product	Policy Agent must be authorized to a security product profile. See z/OS Communications Server: IP Configuration Guide for details about setting up the proper authorization.	EZZ8434I message, along with messages in the log file indicating that the user is not authorized to start Policy Agent and showing an exit code value of 18.
When the SERVAUTH class is active, either: • INITSTACK security product profile is not defined • Policy Agent not permitted to the INITSTACK security product profile.	See "Common AT-TLS startup errors" on page 707 for how to handle this startup error.	EZZ4248E message, and in the Policy Agent log file you see the SYSERR message that indicates that the socket could not open with errno= Resource temporarily unavailable., errno2=74610296.
Unable to read configuration file	• The correct configuration file must be specified. See z/OS Communications Server: IP Configuration Guide for the search order used to locate the main configuration file. • The file must exist. • The permission bits must be correctly set for a z/OS UNIX file system file. • Because multiple configuration files might be configured, you might need to check these files also.	EZZ8434I message, along with messages in the log file indicating that the configuration file could not be opened and showing an exit code value of 1.

Table 48. Common Policy Agent initialization problems (continued)

Problem	Cause/action	Symptom
Unable to load one or more dynamic link libraries (DLLs) needed when Policy Agent is configured as a policy client	Policy Agent must have access to several DLLs at run time when configured as a policy client. These are needed to access PAPI functions and to establish an SSL connection to the policy server. Policy Agent accesses the DLLs using the LIBPATH environment variable. Check that the LIBPATH environment variable is specified, and that it contains the directory in which the DLLs reside. This is normally /usr/lib.	EZZ8780I message, along with messages in the log file indicating problems loading one or more of the following DLLs: • papi.dll • GSKSSL • GSKCMS31

Policy definition problems

If you do not see the expected results when defining policies, use the **pasearch** command to display policies (active or inactive) known by Policy Agent. Use this command to check whether policies are active or inactive and whether they contain the specifications that were expected.

Guidelines:

- Policy rules with complex conditions (using CNF/DNF logic) are processed by Policy Agent to arrive at a "working" set of conditions. These are the only conditions displayed by default using pasearch (use the -o option to display the original set of conditions as specified).
- The pasearch output displays overall time ranges and time of day ranges in UTC format, as well as the specified time zone, if other than UTC.

You can dynamically refresh Policy Agent so that it can pick up any changes made, including changes to policies on the LDAP server (or configuration file). Use the MODIFY procname, REFRESH command to restart Policy Agent from the beginning of its configuration files, or MODIFY procname, UPDATE command to re-read the configuration files.

To check whether QoS policies are being installed and used correctly, use the NETSTAT commands. Use the NETSTAT SLAP or **netstat -j** command to display active QoS policy statistics for QoS policies installed in the stack, as opposed to the policies in Policy Agent. The NETSTAT ALL or **netstat -A** command shows which QoS policy rule (if any) is mapped to active connections.

For further diagnosis of the following policy types, see the topics listed below:

- Intrusion Detection Services (IDS) policy definition problems

  See Chapter 27, “Diagnosing intrusion detection problems,” on page 699 for more information about diagnosing IDS policy definition problems.
- IPSec policy definition problems

  See Chapter 30, “Diagnosing IP security and defensive filter problems,” on page 733 for more information about diagnosing IPSec policy definition problems.
- Application Transparent Transport Layer Security (AT-TLS) policy definition problems

See Chapter 28, "Diagnosing Application Transparent Transport Layer Security (AT-TLS)," on page 707 for more information about diagnosing AT-TLS policy definition problems.

You might encounter some of the policy definition problems listed in Table 49.

Table 49. Policy definition problems

Problem	Cause/action	Symptom
GSKCMS31 DLL not found	Policy Agent must have access to the GSKCMS31 DLL at run time. This is needed for IPSec KeyExchange policies. The IPSec policy being validated failed. Policy Agent accesses the GSKCMS31 DLL using the LIBPATH environment variable. Check that the LIBPATH environment variable is specified, and that it contains the directory in which the GSKCMS31 DLL resides. This is normally /usr/lib.	Policy Agent logs a system error message and object error message.
GSKSSL DLL not found	Policy Agent must have access to the GSKSSL DLL at run time. This is needed for AT-TLS policies. Policy Agent loads the AT-TLS policies into the TCP/IP stack, but because Policy Agent was unable to verify with System SSL that the configured cipher suites were valid, they are validated when the TLS/SSL environment is initialized for TCP/IP connections. If any values are not valid within the cipher suites, this could result in TCP/IP connections failing. Policy Agent accesses the GSKSSL DLL using the LIBPATH environment variable. Check that the LIBPATH environment variable is specified, and that it contains the directory in which the GSKSSL DLL resides. This is normally /usr/lib.	Policy Agent logs a system error message and warning message.
Version 1 QoS policies to version 2 QoS policies conversion Semantic differences exist between version 1 and version 2 policy definitions. **Restriction:** Currently only version 2 semantics are supported. When the policies are processed by Policy Agent, version 1 policy semantics are converted to version 2 semantics. See Note 1.	The following circumstances might lead to problems: • When converting such policies to version 2, be sure to also swap the source and destination attributes when the version 1 Direction is Inbound. The specified interface is also related to Direction. In version 1 only a single interface is specified, while both inbound and outbound interfaces are specified in version 2. When migrating a version 1 policy, be sure to specify an InboundInterface for Direction Inbound, and an OutboundInterface for Direction Outbound. • When converting version 1 rules with Direction Both specified, create two version 2 rules, one for each direction. Also, specify InboundInterface for the inbound rule and OutboundInterface for the outbound rule, if the version 1 rule specified both Interface and Direction Both. • When converting policies with different PolicyScope values, be sure to logically merge the scopes in the version 2 policy action. Any such merge should always result in a PolicyScope value of Both.	Discrepancies between version 1 and version 2 policy definitions.

Table 49. Policy definition problems (continued)

Problem	Cause/action	Symptom
Policy groups or rules are discarded when defined on an LDAP server.	Policy groups and policy rules defined on an LDAP server can refer to other LDAP objects (such as policy actions or time periods). When any referenced object cannot be found on the LDAP server, the referencing object is discarded. Specify the correct reference Distinguished Names on LDAP objects that reference other objects.	Discarded policy groups or rules
Policies with complex conditions (using CNF or DNF) are not mapping correctly.	Because some conditions are logically ANDed, a result that is not valid can occur. For example, two or more distinct interfaces cannot be ANDed and still be true. Or two non-overlapping port ranges also cannot be ANDed. Policy Agent tries to detect these types of errors and discard the policy rules with an error message, but there are cases that cannot be detected (for example, logical ANDs between CNF/DNF levels, or when negated conditions are used). In these cases, a policy rule can be installed that can never be true. Similar problems could occur when ORing conditions. For example, a very broad condition might map much more traffic than was intended, simply because it is one of a set of conditions that is ORed together. Use the **pasearch** command to display policy rules with complex conditions. By default, the "working" set of conditions is displayed (after Policy Agent has attempted to collapse and summarize the complex conditions). This working set includes the summary of each condition level, as well as the overall "global" summary condition. Use the pasearch -o option to also display the original set of specified conditions. This helps to show how the working set was derived.	Difficultly configuring complex policy conditions using CNF or DNF.
Wrong policy being mapped to traffic	At times, two or more policy rules are logically mapped to the same set of traffic packets. When this happens, the rule with the highest weight is selected. The weight depends on two factors. When the policy rule priority is not specified, the weight depends on the number of attributes specified in the policy conditions. When policy rule priority is specified, the weight is the specified priority plus 100, which is always higher than the weight derived from counting the number of attributes. If more than one rule is found with the same weight, the first such rule is selected to be mapped. Be sure to specify priority in policy rules to better control situations where multiple rules map to the same set of traffic.	Policy rule priority settings are inadequate to control situations where multiple rules map to the same set of traffic.

Table 49. Policy definition problems (continued)

Problem	Cause/action	Symptom
Policies are not installed in the TCP/IP stack.	Perform the following actions, based on what caused the problem: • The stack in which policies should be installed must be configured using a TcpImage statement in the Policy Agent configuration file. • The time periods configured in the policies must be correct. Verify the specifications of the day of week and time of day are correct. Verify that the specified time zone is correct. For time zones other than local time, the specified time periods might not be currently active. • If the stack was started or restarted after Policy Agent was started, check that the temporary file (/tmp/*tcpname*.Pagent.tmp) used by the stack to inform Policy Agent of restarts has not been deleted. Perform the following steps to diagnose QoS problems: • Issue **pasearch -q** to see all QoS policies that are active in Policy Agent. See z/OS Communications Server: IP System Administrator's Commands for more information about the **pasearch -q** command. If you are running multiple stacks, ensure that pasearch is reporting on the stack you are interested in. • Issue NETSTAT SLAP or **netstat -j** command to see how the stack mapped your QoS statement. See z/OS Communications Server: IP System Administrator's Commands for more information about the **netstat** command. If you are running multiple stacks, ensure that your resolver configuration correctly identifies the stack you are interested in. Ensure that your QoS policies are correctly defined. For more information about policy-based networking, see z/OS Communications Server: IP Configuration Guide. • See Chapter 27, "Diagnosing intrusion detection problems," on page 699 for more information about diagnosing IDS policy definition problems. • See Chapter 30, "Diagnosing IP security and defensive filter problems," on page 733 for more information about diagnosing IP security problems. • See Chapter 28, "Diagnosing Application Transparent Transport Layer Security (AT-TLS)," on page 707 for more information about diagnosing AT-TLS policy definition problems.	Unexpected or missing set of policies.

Table 49. Policy definition problems (continued)

Problem	Cause/action	Symptom
QoS policies not mapping to the expected traffic	Incorrectly specified selection criteria on the PolicyRule statement for the policy. If you think data traffic should be mapped to certain policies, but is not, check to make sure you have specified the selection criteria correctly on the PolicyRule statement for the policy. For example, TCP policies are mapped on a per connection basis, whereas for UDP and RAW, the policy is mapped on a per packet basis. As an example of TCP traffic, consider an ftp GET request from a remote client. The connection request from the client is mapped as inbound data, while the data flow is mapped as outbound data. You can use either source or destination fields in the policy rule to map both traffic flows, but the definitions must be consistent with this way of mapping. Check that the policy is not unnecessarily restrictive in its specification of IP addresses and ports. For RSVP scoped policies, remember that the policy is only mapped to data traffic while an RSVP reservation is in effect.	A blank policy rule name is displayed for an active connection using the NETSTAT ALL or **netstat -A** command.
Timing windows when switching policies based on time	If policy rules are defined such that different sets of policies are activated at different times (for example at each shift), be aware of nonoverlapping vs. overlapping time specifications. For example, if Rule1 is active from 00:00 to 07:29, and Rule2 is active from 07:30 to 04:00, there is a one minute interval gap between these 2 rules. Because the minimum time resolution used by Policy Agent is one minute, there is a period of one minute when neither policy is active.	Different sets of policies are activated at different times (for example at each shift).
Policies defined in an MVS data set are not being installed.	When an MVS data set is used to define policies, ensure that sequence numbers are not part of the file, because these cause parsing errors. In ISPF, use the NUMBER OFF and UNNUM or NUMBER OFF or UNNUM commands to remove the sequence numbers.	Parsing errors occur.

Note 1. Be aware of the following processing behavior:

- In version 1, source always meant local, while destination always meant remote. In version 2, source and destination mean exactly what they imply. When version 1 policies specify Direction Inbound, the semantics for source and destination are opposite between the two versions. As a result, although the specified source and destination attributes are displayed as they are specified by the **pasearch** command, the attributes are swapped when the policies are installed in the stack.
- Similarly, when Direction Both is specified in a version 1 policy, the following policies are installed in the stack:
  - Outbound direction with source and destination attributes intact
  - Inbound direction with the attributes swapped
- PolicyScope values exist in both the policy rule and action in version 1, but only in the policy action in version 2. For any policies that specified different PolicyScope values for the rule and the associated action in version 1, the scope values are merged in the policy action. For example, if the rule specified PolicyScope Both, and the associated action specified PolicyScope DataTraffic, the resulting scope value in the policy action is Both.

Policy client connection problems

When acting as a policy client, Policy Agent needs to connect to a policy server. The policy client can be configured with just a primary, or both a primary and a backup, policy server. See z/OS Communications Server: IP Configuration Guide, Policy Agent and policy applications for more information about how the policy client connects to a policy server.

If the policy client does not connect successfully, run Policy Agent on the policy client and policy server with the -d 128 startup option, and check the log files for error conditions. Connection problems are indicated by message EZZ8780I or message EZZ8782I. Check the log files for the specific error encountered.

Table 50 describes common policy client connection problems.

Table 50. Common policy client connection problems

Problem	Cause/action	Symptom
Incorrect configuration on the policy client or policy server	• The policy server must be configured with the ClientConnection statement specifying the port to which policy clients connect. • If you use secure connections from any policy clients, the policy server must be configured with AT-TLS policies that allow those policy clients to establish SSL connections to the policy server. • The policy client must be configured with the ServerConnection statement specifying the host name or IP address, and port of the primary and optional backup policy server, as well as connection retry information. • If you want to use a secure connection to the policy server, you must configure the policy client with SSL information about the ServerConnection statement. See the policy-based networking topic in z/OS Communications Server: IP Configuration Guide for details about setting up the correct configuration.	Message EZZ8780I or EZZ8782I, along with messages in the log files indicating the particular connection problem details.

Table 50. Common policy client connection problems (continued)

Problem	Cause/action	Symptom
Incorrect SSL configuration on the policy client or policy server	If you use secure connections from the policy client: • The policy server must be configured with AT-TLS policies that allow the policy clients to establish SSL connections to the policy server. • The policy server must be configured with a certificate that allows the policy clients to authenticate the server. • If a self-signed server certificate is used, the policy client must import the server's certificate into the client's key ring. • The ServerConnection statement on the policy client must be configured with the correct SSL parameters. See the policy-based networking topic in z/OS Communications Server: IP Configuration Guide for details about setting up the correct configuration.	Message EZZ8780I or EZZ8782I, along with messages in the log files indicating the particular connection problem details.
Mismatched security configuration between the policy client and policy server	The configuration on the policy client must match the configuration on the policy server with respect to SSL and AT-TLS: • If the policy client is configured with SSL parameters on the ServerConnection statement, the policy server must have an AT-TLS policy that protects connections from that policy client. • If the policy client is not configured with SSL parameters on the ServerConnection statement, the policy server must not have an AT-TLS policy that protects connections from that policy client. Use the **pasearch** command to display the AT-TLS policies on the policy server, and verify that the selection criteria in the policy rules select only those policy clients that use SSL. Look for policy rules that specify the port specified on the ClientConnection statement as the local port, and in particular, verify that the remote IP address and remote port parameters on those policy rules are correct for your configuration. See the policy-based networking topic in z/OS Communications Server: IP Configuration Guide for details about setting up the correct configuration.	Message EZZ8780I or EZZ8782I, along with messages in the log files indicating the particular connection problem details.

Table 50. Common policy client connection problems (continued)

Problem	Cause/action	Symptom
Incorrect certificate name specified on the ServerSSLName parameter on the ServerConnection statement	• If the AT-TLS policy on the policy server specifies HandshakeRole Server, the ServerSSLName parameter on the ServerConnection statement on the policy client must specify the name of the server's certificate. • If the AT-TLS policy on the policy server specifies HandshakeRole ServerWithClientAuth, the ServerSSLName parameter on the ServerConnection statement on the policy client must specify the name of the client's certificate. See the policy-based networking topic in z/OS Communications Server: IP Configuration Guide for details about setting up the correct configuration.	Message EZZ8780I or EZZ8782I, along with messages in the log files indicating the particular connection problem.
Policy client not authorized to access policy server	• The policy server must be configured with one or more user IDs and credentials for the set of policy clients that are authorized to connect. **Rule:** If you use a password for credentials, the password must match the password configured using the AuthBy password parameter on the PolicyServer statement on the policy client. • The policy client must be configured with a PolicyServer statement for each stack that will retrieve policies from the policy server, indicating the user ID and credentials that will be used to access the policy server. See the policy-based networking topic in z/OS Communications Server: IP Configuration Guide for details about setting up the correct authorization.	Message EZZ8780I or EZZ8782I, along with messages in the log files indicating the particular authorization problem details.
Incorrect passticket configuration on the policy client or policy server	If the policy client is configured to use a passticket on the PolicyServer statement, the proper PTKTDATA class profiles must be defined on both the policy server and policy client. See the policy-based networking topic in z/OS Communications Server: IP Configuration Guide for details about setting up the correct configuration.	Message EZZ8780I or EZZ8782I, along with messages in the log files indicating the particular connection problem details.

Table 50. Common policy client connection problems (continued)

Problem	Cause/action	Symptom
The policy server is not listening on the port defined on the ClientConnection statement.	If the ClientConnection statement is configured on the policy server, the port specified on this statement may need to be reserved using the PORT statement in the TCP/IP profile. See z/OS Communications Server: IP Configuration Guide for details about setting up the correct configuration.	Message EZZ8788I, along with messages in the log files, indicating the particular connection problem details.
Duplicate policy client name reported	If you use the configuration file import service on the policy server, you might encounter a duplicate policy client name for a policy client. The reason for this is that temporary names are generated in order to process a configuration file import. If a policy client tries to connect to the policy server while a configuration file import is in progress, it's possible that the policy client name matches the generated temporary name. If this happens, issue a MODIFY UPDATE command on the policy client to cause it to reconnect to the policy server, once the configuration file import service has completed.	Message EZZ8781I followed by message EZZ8782I, along with messages in the log files indicating a duplicate policy client name was detected.

Policy client retrieval problems

When acting as a policy client, Policy Agent retrieves policies for one or more policy types, on behalf of one or more stacks, from a policy server. The choice of local or remote policy retrieval can be made separately for each policy type, and for each configured stack. See z/OS Communications Server: IP Configuration Guide, Policy Agent and policy applications for more information about policy client retrieval of remote policies.

If the policy client does not successfully retrieve policies, run Policy Agent on the policy client and policy server with the -d 128 startup option, and check the log files for error conditions. Retrieval problems are indicated by message EZZ8438I. Check the log files for the specific error encountered.

Table 51 on page 672 describes common policy client retrieval problems.

Table 51. Common policy client retrieval problems

Problem	Cause/action	Symptom
Incorrect configuration on the policy client or policy server	• The policy server should be configured with one or more DynamicConfigPolicyLoad statements that match the client name. The DynamicConfigPolicyLoad statement determines the configuration files that get loaded after a policy client successfully connects. If a matching DynamicConfigPolicyLoad statement is not found, the policy server will attempt to load policies from a default file. Ensure that the correct set of DynamicConfigPolicyLoad statements is specified, and that the correct configuration files are specified on these statements. • The policy client must be configured with a PolicyServer statement for each stack that will retrieve policies from the policy server. The ClientName specified on this statement is used to match a DynamicConfigPolicyLoad statement on the policy server. If the ClientName parameter is not specified, the default client name used is *remotesysname_tcpimage* where: *remotesysname* value is the policy client system name and *tcpimage* value is the policy client image name See z/OS Communications Server: IP Configuration Guide , Policy Agent and policy applications, and z/OS Communications Server: IP Configuration Reference, general configuration file statements topic for more information.	Incorrect or no policies retrieved from the policy server.

Table 51. Common policy client retrieval problems (continued)

Problem	Cause/action	Symptom
Incorrect regular expressions coded on the DynamicConfigPolicyLoad statement	The DynamicConfigPolicyLoad statements can be configured with regular expressions to match against policy client names. Regular expressions are very powerful, but also can be complex, and might not produce results that are intuitive. For example, the expression [a-z] matches any lower case alphabetic character, which means that any string containing at least one such character will match. As another example, the expression [^abc] means any character except a, b, or c matches. So the only strings that won't match are those containing ONLY the characters a, b, or c. See z/OS Communications Server: IP Configuration Guide , Policy Agent and policy applications, and z/OS Communications Server: IP Configuration Reference, DynamicConfigPolicyLoad statement for more information.	Incorrect or no policies retrieved from the policy server.
Policy client not authorized to access policies on the policy server	The policy server must be configured with SERVAUTH profiles that allow the policy clients to access policies. The format of the SERVAUTH profiles is: EZB.PAGENT.*sysname.image.ptype* where: • *sysname* is the policy server system name • *image* is the policy client name **Rule:** The *image* portion of the profile name on the policy server must match or include the name of the policy clients. Each policy client name is configured or defaulted using the ClientName parameter on the PolicyServer statement. • *ptype* is the policy type (QOS, IDS, IPSEC, ROUTING, or TTLS) See z/OS Communications Server: IP Configuration Guide, Policy Agent and policy applications for more information.	Incorrect or no policies retrieved from the policy server.

Import requestor connection problems

The import requestor connects to a Policy Agent. The IBM Configuration Assistant for z/OS Communications Server can be an import requestor. See z/OS Communications Server: IP Configuration Guide for information about the configuration file import service.

If the import requestor does not connect successfully, run Policy Agent with the -d 128 startup option, and check the log files for error conditions. Connection problems are indicated by message EZD1578I in the log. Check the log files for the specific error encountered.

Table 52 describes common import requestor connection problems.

Table 52. Common import requestor connection problems

Problem	Cause/action	Symptom
Incorrect configuration on Policy Agent	You must configure the Policy Agent with the ServicesConnection statement specifying the port and TCP/IP stack name to which the import requestor will connect. See z/OS Communications Server: IP Configuration Guide , configuration file import service, for details about setting up the correct configuration.	Message EZD1578I, along with messages in the log files, indicating the particular connection problem details.

Table 52. Common import requestor connection problems (continued)

Problem	Cause/action	Symptom
Incorrect SSL configuration on the Policy Agent or import requestor	• If the import requestor is using SSL, you must configure the Policy Agent with a SAF keyring and the Security parameter set to Secure on the ServicesConnection statement. – The Policy Agent generates and installs an AT-TLS policy that allows the import requestors to establish SSL connections to the Policy Agent. – You must configure a certificate in a SAF keyring that allows the import requestors to authenticate the server. **Tip:** This option only supports TLSv1.0 and is not recommended for securing the import connection. The recommended approach is to specify Security Basic and supply user defined AT-TLS policies that provide the required protection. • If the import requester is not using SSL, you must configure the Policy Agent with the Security parameter set to Basic on the ServicesConnection statement and use a default unsecure connection. You must not configure an AT-TLS policy that includes the port configured on the ServicesConnection statement. • If the import requester is using user defined SSL, you must configure the Policy Agent with the Security parameter set to Basic on the ServicesConnection statement. – You must define AT-TLS policies for the ServicesConnection ImageName and port that are configured for this secure SSL connection. – You must configure the matching import requester and Policy Agent AT-TLS policies. See z/OS Communications Server: IP Configuration Guide , configuration file import service, for details about setting up the correct configuration.	Message EZD1578I, along with messages in the log files, indicating the particular connection problem details.

Table 52. Common import requestor connection problems (continued)

Problem	Cause/action	Symptom
The Policy Agent did not issue message EZD1576I indicating it is ready for services connection requests.	If you are using secured connections for import requestors, and have AT-TLS policies configured on a policy server, the Policy Agent waits for the remote AT-TLS policies to be retrieved and installed before installing the generated AT-TLS policy for the port specified on the ServicesConnection statement. If the policy server is down or cannot be contacted immediately, the generated AT-TLS policy cannot be installed and the Policy Agent does not listen for import requestor connections. • Verify the policy server is available and the Policy Agent is active. • You might consider using a backup policy server to handle policy client connections when the primary is not available. • The MODIFY SRVLSTN command could be used to force the generated AT-TLS policy to be installed before the remote AT-TLS policies are installed. • Run the policy client and policy server with debug level 128 and check the Policy Agent log files to determine the cause of any connectivity problems. See z/OS Communications Server: IP System Administrator's Commands for information about the MODIFY command and z/OS Communications Server: IP Configuration Guide, for AT-TLS data protection information.	Message EZD1576I is not issued, and import requestors cannot connect to the Policy Agent. Message EZD1578I, along with messages in the log files, indicating the particular connection problem details.
The Policy Agent did not issue message EZD1576I indicating it is ready for services connection requests.	If you are using secured connections for import requestors, and have local or remote AT-TLS policies configured that contain errors, Policy Agent waits for the local or remote AT-TLS policies to be installed. • Correct the configured AT-TLS policies and refresh policies. • The MODIFY SRVLSTN command could be used to force the generated AT-TLS policy to be installed before the local or remote AT-TLS policies are installed. See z/OS Communications Server: IP System Administrator's Commands for information about the MODIFY command and z/OS Communications Server: IP Configuration Guide, for AT-TLS data protection information.	Message EZZ8438I, indicating errors in the local or remote AT-TLS policies for a TCP/IP image, where the secured connections for import requestor is requested. Message EZD1578I, along with messages in the log files indicating the particular connection problem details. Message EZD1576I is not issued, and import requestors cannot connect to the Policy Agent.

Table 52. Common import requestor connection problems (continued)

Problem	Cause/action	Symptom
Import requestor does not successfully connect to the Policy Agent.	If you are using secured connections for import requestors and the SAF keyring is correct but the connection from the import requestor fails, (indicating key ring problems) check the following: • If the key ring certificate has expired, then update the expiration date and issue the MODIFY SRVLSTN command for Policy Agent to reinstall the generated AT-TLS policy and to restart the listen for services requestor connections • If the contents of the key ring has changed, but the key ring name is unchanged, issue the MODIFY SRVLSTN for Policy Agent to reinstall the generated AT-TLS policy and to restart the listen for services requestor connections See z/OS Communications Server: IP System Administrator's Commands for information about the MODIFY command and z/OS Communications Server: IP Configuration Guide, for AT-TLS data protection information.	Message EZD1576I is issued, but import requestors cannot connect to the Policy Agent.
The Policy Agent is not listening on the port defined on the ServicesConnection statement.	If the ServicesConnection statement is configured, the port specified on this statement may need to be reserved using the PORT statement in the TCP/IP profile. See z/OS Communications Server: IP Configuration Guide for details on setting up the correct configuration.	Message EZD1578I, along with messages in the log files, indicating the particular connection problem details.
Import requestor not authorized to access Policy Agent system	The Policy Agent system must be configured with one or more user IDs and credentials for the set of import requestors that are authorized to connect. **Rule:** If you use a password for credentials, the password must match the password configured on the import requestor. If you use the IBM Configuration Assistant for z/OS Communications Server as the import requestor, the user ID and password are configured on the Import Policy Data panel or request panels for discovery import (for example Discover Stack Local Addresses panel). See z/OS Communications Server: IP Configuration Guide for details on setting up the correct configuration.	Message EZD1578I, along with messages in the log files, indicating the particular connection problem details.

Import requestor retrieval problems

The import requestor retrieves import policies from the Policy Agent.

Table 53 describes common import requestor retrieval problems.

Table 53. Common import requestor retrieval problems

Problem	Cause/action	Symptom
Import requestor not authorized to access import policies on the Policy Agent.	The Policy Agent must be configured with SERVAUTH profiles that allow the import requestor to access import policies. The format of the SERVAUTH profiles is: `EZB.PAGENT.`*`sysname.image.ptype`* where: • *sysname* is the local system name • *image* is the import request name **Rule:** The *image* portion of the profile name on the policy server must match or include the import request name. If you use the IBM Configuration Assistant for z/OS Communications Server as the import requestor, the import request name is configured on the Flat File Import panel. • *ptype* is the type (CFGSERV, IDS, IPSEC, ROUTING, or TTLS) See z/OS Communications Server: IP Configuration Guide, configuration file import service, for details about setting up the correct configuration.	Incorrect or no import policies retrieved from the Policy Agent.

LDAP object retrieval problems

Before you begin, if you are having problems receiving policies from an LDAP server, run Policy Agent with the -d 1 or 2 startup options.

In Table 54, select actions as indicated according to the problem you are experiencing.

Table 54. LDAP object retrieval problems

Problem	Cause/action	Symptom
Unable to connect to the LDAP server	Check the attributes specified on the ReadFromDirectory statement in the configuration file that relate to the LDAP server connection. These include the primary and backup server addresses and ports, the user ID and password, and SSL parameters.	Message EZZ8440I is issued to the console. If Policy Agent fails to connect to the LDAP server, check the log file for the specific error encountered. The Policy Agent keeps trying to connect to the server, using a sliding time window (one minute, then at five minute intervals, with the maximum time between connect attempts being 30 minutes). **Tip:** If a backup LDAP server is configured, the EZZ8440I message is only issued if neither the primary or backup server can be connected.

Table 54. LDAP object retrieval problems (continued)

Problem	Cause/action	Symptom
No objects, or incorrect objects, retrieved from the LDAP server	Check that the schema version specified on the ReadFromDirectory statement in the configuration file matches the version defined on the LDAP server. The different versions are distinguished by the set of supported object classes. See z/OS Communications Server: IP Configuration Guide for supported schema object classes.	Missing or incorrect policies are displayed by the **pasearch** command, or the NETSTAT SLAP or **netstat -j** commands.
Wrong set of objects retrieved from the LDAP server	Check that the search and selection criteria specified on the ReadFromDirectory statement in the configuration file are correct. For version 1 policies, verify that the correct Base and SelectedTag attributes are used. For version 2 and later policies, check the SearchPolicyBaseDN, SearchPolicyGroupKeyword, SearchPolicyKeyword, and SearchPolicyRuleKeyword attributes.	Missing or incorrect policies are displayed by the **pasearch** command, or the NETSTAT SLAP or **netstat -j** commands.
LDAP DLL not found **Restriction:** Policy Agent must have access to the LDAP DLL at run time.	Check that the LIBPATH environment variable is specified, and that it contains the directory in which the LDAP DLL (GLDCLDAP) resides. This is normally /usr/lib. Policy Agent accesses the LDAP DLL using the LIBPATH environment variable.	Policy Agent terminates unexpectedly with a CEEDUMP. The reason for termination in the CEEDUMP indicates that the LDAP DLL (GLDCLDAP) was not found.
Version 1 policies not shared among multiple TCP/IP stacks	Policy Agent uses two attributes when it searches an LDAP server for version 1 policies that apply to a given TCP/IP image. One attribute is the TCP/IP image name and the other is a selector tag. The selector tag attribute can be defined such that LDAP scopes the search. The TCP/IP image name attribute is set by default to scope the search for a particular image. Each of the two attributes (TCPImageName and SelectorTag) is a multivalue field, meaning you can specify TCPImpageName/SelectorTag multiple times in one object defined to LDAP. Both multiple MVS images and multiple TCP/IP stacks can exist. If a policy object is to be used in multiple MVS LPARs, that object can have multiple SelectorTag attributes defined, one for each LPAR. If a policy object is to be used in multiple TCP/IP images, that object can have multiple TCPImageName attributes defined, one for each image.	Version 1 policies not shared among multiple TCP/IP stacks

LDAP object storage problems

Policies can be defined on an LDAP server using the appropriate definitions, known as schemas. The policies are defined as object classes with certain attributes, which are a superset of the attributes that can be defined in a local file using the PolicyAction and PolicyRule statements. Policy Agent acts as an LDAP client to communicate with and retrieve policies from an LDAP server. Policy Agent uses an LDAP DLL to perform its LDAP client functions.

Before you begin, if you are having problems initializing the LDAP server with the Policy Agent schema definitions or adding policy objects to the server, perform the following steps to diagnose LDAP object storage problems.

In Table 55, select actions as indicated according to the problem you are experiencing.

Table 55. LDAP object storage problems

Problem	Cause/action	Symptom
Unable to add the Policy Agent schema definitions to an LDAPv3 server	The Policy Agent LDAPv3 schema definition files are shipped as the following sample files: • pagent_r8qosschema.ldif • pagent_r5idsschema.ldif These files need to be installed on the LDAP server in the proper order as an object in the server's database, rather than as configuration information. This process is known as schema publication. See RFCs 1804 and 2251. The files need to be specified on **ldapmodify** commands to modify the cn:schema entry in the server's database, in the order as specified in z/OS Communications Server: IP Configuration Guide. Verify that the <suffix> value on the first noncomment line of these files has been changed to the suffix value defined for your LDAP server, as explained in the prologues in these files. For more information about installing the schema definition files, see z/OS Communications Server: IP Configuration Guide.	Symptoms can include error messages issued by the server. Because server implementations are different, check the documentation for your server for the types and locations of error or log messages.

Table 55. LDAP object storage problems (continued)

Problem	Cause/action	Symptom
Unable to add policy objects to an LDAP server	Check the following: 1. Are the Policy Agent schema definitions installed on the LDAP server? 2. Are the correct object classes identified for any attributes you have defined in the object? For example, the ibm-policySubtreesAuxContainedSet attribute is defined for the ibm-policySubtreesPtrAuxClass object class. 3. Does the server recognize all of your objects?	Symptoms can include error messages issued by the server. Since server implementations are different, check the documentation for your server for the types and locations of error or log messages. A typical error message might indicate object class violation. There are several possible reasons for an LDAP server rejecting a policy object. The following symptoms correspond to the numbered actions in the cause and action column. 1. If the server does not know about policy attributes or object classes, then it fails any objects that contain them. 2. If you define a policy object with this attribute attached, but do not include the object class value, the server flags the object as an object class violation. 3. The symptoms for this are missing objects when you search the server or errors when adding the objects. Some servers can impose strict syntax rules on ldif files that contain objects. • Lines that separate objects might need just a single newline character. If the separator lines contain other characters, the following object is processed as a continuation of the previous object. If the object file was transferred using FTP from a host, character translation might result in characters other than newlines separating objects. These additional characters must be removed. • There must be no blanks at the ends of lines.

Policy Agent and Sysplex distribution problems

The Policy Agent sysplex distributor (SD) performance monitor function can be used to calculate outbound network performance information, such as TCP packet loss and timeout ratios, for applications being distributed to on SD target nodes. The calculated performance information is in the form of QoS weight fractions calculated for each DVIPA/Port service level. The QoS weight fractions are used to adjust the WLM weight: the higher the Qos fraction calculation, the lower the adjusted WLM weight. For more information about QoS fractions, see the section

about Sysplex Distributor Policy Performance Monitoring Configuration in the z/OS Communications Server: IP Configuration Guide.

Steps for diagnosing Policy Agent/Sysplex distribution problems

Diagnose Policy Agent/Sysplex distribution problems.

Before you begin

If you suspect problems with the calculated QoS weight fractions, run Policy Agent with debug level 4 or 8.

Procedure

Perform the appropriate steps for the problem you are experiencing:

- For debug level 4, Policy Agent displays a summary calculation for each DVIPA/Port XCF address and service level. The summary information includes the retransmit fraction, connection limit fraction, throughput fraction, and the final QoS fraction that resulted. For example:

  ```
  Calculating for DVIPA: 193.1.1.36, Port: 8000, XCF@: 193.1.1.36, SLName:
   'Gold_Service'
   Fractions: rexmit: 0, connLimit: 100, thruput: 0 QoS used: 100
  ```

- For debug level 8, Policy Agent displays the intermediate values used to generate the above fractions. For example:

  ```
  Calculating for DVIPA: 193.1.1.36, Port: 8000, XCF@: 193.1.1.36, SLName:
   'Gold_Service'
   Retransmit Fraction: 0 (Retransmit Bytes: 544, Timeouts: 1,
   Octets Sent: 81362424, Segments Sent: 143194)
   Connection limit Fraction: 100 (Max Connections: 3, Active Connections: 3)
   Throughput Fraction: 0 (Out Bytes: 81362424, Throughput: 10848,
   Conn Throughput: 3616 Profile Rate: 0, Min Rate 2000)
   QoS Fraction used : 100
  ```

  Guideline: If the throughput fraction gets set to 100% for any service level, message EZZ8447I is issued. To see which service levels caused this message to be issued, run Policy Agent with debug level 8 and check the log file.

What to do next

For more information see Chapter 11, "Diagnosing dynamic VIPA and sysplex problems," on page 383.

Memory allocation/leakage problems

Policy Agent allocates memory for many resources, such as:

- Policy rules and actions
- Sysplex Distributor lists and weight fraction arrays
- Policy performance data arrays
- LDAP search results

If it appears that Policy Agent is using too much memory, or memory leakage is suspected, use the following tools, possibly in conjunction with other tools outside the scope of Policy Agent, such as dump formatters and Language Environment memory tracing.

Use the -m startup option to keep track of all Policy Agent memory allocation and free requests. All memory allocations are recorded in a memory trace buffer, and

all memory free requests find the corresponding entry and remove it. If this option is specified, Policy Agent automatically reports any memory leakage at termination time, because any entries left in the buffer after all memory free requests have been processed are by definition memory leaks. Note that if the memory trace buffer fills up, the memory trace function is dynamically turned off and no more memory tracing is performed. If this occurs, specify a larger value for the -m startup option when Policy Agent is restarted.

Use the MODIFY MEMTRC command to log a snapshot of Policy Agent memory allocations. This command dumps the contents of the memory trace buffer to the log file. As a result, it only has an effect when the -m startup option was specified.

Use debug level 16 to record memory allocation and free requests inline in the log file. This debug level is independent of the -m startup option. Note that using this debug level can result in significantly more information being recorded, so specify larger and/or more log files using the PAGENT_LOG_FILE_CONTROL environment variable.

Chapter 26. Diagnosing RSVP agent problems

The z/OS UNIX RSVP Agent provides end-to-end resource reservation services on behalf of applications. This topic provides information and guidance to diagnose z/OS UNIX RSVP Agent problems and contains the following sections:

- "Overview"
- "Policies and RSVP processing" on page 687
- "Gathering diagnostic information" on page 688
- "Diagnosing RSVP agent problems" on page 688

Overview

The RSVP Agent provides an RSVP application programming interface (RAPI) for QoS-aware applications to use. Applications use RAPI to register their intent to use RSVP services, to describe their data traffic, and to explicitly request that network resources be reserved on their behalf. The RSVP Agent communicates with its peers (other RSVP Agents running on z/OS or other platforms) in the network, with QoS-aware sender and receiver applications, and with the TCP/IP stack to affect resource reservations. See RFC 2205 for more information about RSVP, and z/OS Communications Server: IP Programmer's Guide and Reference for more information about RAPI.

The following terms must be defined to understand RSVP processing:

Quality of Service (QoS)
: The overall service that a user or application receives from a network, in terms of throughput, delay, and such.

QoS-Aware Application
: An application that explicitly requests QoS services from the RSVP agent.

Service Differentiation
: The ability of a network to provide different levels of QoS to different users or applications based on their needs.

Service Level Agreement (SLA)
: A contract in business terms provided by a network service provider that details the QoS that users or applications are expected to receive.

Service Policy
: Administrative controls for a network, in order to achieve the QoS promised by a given SLA.

Integrated Services
: A type of service that provides end-to-end QoS to an application, using the methodology of resource reservation along the data path from a receiver to a sender.

Differentiated Services
: A type of service that provides QoS to broad classes of traffic or users, for example all FTP traffic to a given subnet.

Resource ReSerVation Protocol (RSVP)
: A protocol that provides for resource reservation in support of Integrated Services.

© Copyright IBM Corp. 2000, 2015

Reservation types, styles, and objects

There are two types of Integrated Services reservations used by the RSVP Agent:

Controlled Load
This reservation type is designed to make the network behave as though it were not loaded, even if one or more of the network elements are experiencing a heavy traffic load. See RFC 2211 for more information about this service.

Guaranteed
This reservation type is designed to allow the network to compute the maximum delay data traffic receives from the network, based on the traffic specification and other known data. See RFC 2212 for more information about this service.

In addition, there are three styles of reservation, depending on how the receiver wants to apply the reservation to its senders:

WF (Wildcard Filter)
This style applies a single reservation request to all senders.

FF (Fixed Filter)
This style pairs a given reservation request to a given sender. In this way, the receiver can apply a different reservation to each of its senders.

SE (Shared Explicit)
This style applies a single reservation to a list of senders. This differs from the WF style in that the list of senders is finite. Additional senders that appear in the future do not automatically inherit an SE style reservation.

Several objects are used in RSVP and RAPI to describe data traffic and reservations. These objects are as follows:

Tspec (traffic specification)
The Tspec is used to describe the sending application data traffic characteristics. It consists of an object known as a token bucket and other related values. A token bucket is a continually sustainable data rate, and the extent to which the rate can exceed the sustainable level for short periods of time. More detail concerning token buckets and other Integrated Services parameters and processing can be found in RFCs 2210, 2211, 2212, and 2215.

The Tspec contains these values:

r Token bucket rate, in bytes per second

b Token bucket depth, in bytes

p Peak rate, in bytes per second

m Minimum policed unit (minimum packet size to be considered), in bytes

M Maximum packet size (MTU), in bytes

Rspec (guaranteed receiver specification)
An Rspec consists of two values that further describe a reservation request when Guaranteed service is being used:

R Requested rate, in bytes per second

S Slack term, in microseconds

Flowspec (reservation specification)

The flowspec is the object used by a receiver application to indicate an actual reservation to be made. The actual makeup of the flowspec depends on the type of reservation. For Controlled Load, the flowspec takes the same form as the sender Tspec (although the form is the same, the receiver might specify different values than the sender). For Guaranteed, the flowspec takes the form of a Tspec followed by an Rspec.

Policies and RSVP processing

Policies can be defined with RSVP scope. The RSVP Agent obtains a service policy for which traffic is mapped (if any) from the Policy Agent when an application that is using RAPI indicates it is a sender (when the Tspec is first provided), or when it requests a reservation as a receiver (when the Rspec is first provided for Guaranteed service). At both of these times, if a service policy is defined that maps to the data traffic, the RSVP Agent uses values in the service policy to limit the request from the application. Specifically, the following are limited:

- Total number of RSVP flows.

  The MaxFlows keyword on the PolicyAction statement of the policy definition can be used to limit the total number of application flows that use RSVP services.
- Tspec token bucket values.

  The MaxRatePerFlow and MaxTokenBucketPerFlow keywords on the PolicyAction statement of the policy definition can be used to limit the r and b values, respectively, in the sender supplied Tspec.
- Rspec values.

  The MaxRatePerFlow keyword on the PolicyAction statement of the policy definition can be used to limit the R value in the receiver supplied Rspec.
- Reservation type.

  The FlowServiceType keyword on the PolicyAction statement of the policy definition can be used to limit the type of reservation requested. A Guaranteed type request is considered to be "greater than" a Controlled Load type request. So if an application requests Guaranteed, but the policy limits the type to Controlled Load, the reservation uses Controlled Load.

RSVP processing proceeds as follows.

When an application uses RAPI to indicate it is a sender, the RSVP Agent packages the sender Tspec (along with other information) in an RSVP PATH packet, and sends the packet to the final destination. The packet is sent using RAW sockets, with the IP Router Alert option set. This option causes each router that supports RSVP to intercept the PATH packet, for the purpose of remembering the PATH request, and to insert a "previous hop" object into the packet, which is then sent again to the final destination. This causes the packet to eventually arrive at the destination, with all RSVP routers in the data path aware of the RSVP flow.

At the destination, the RSVP Agent passes the PATH packet to the application, using RAPI. The receiver application uses the Tspec and other information to arrive at a reservation request (flowspec). The receiver application uses RAPI to pass this flowspec to the RSVP Agent. The RSVP Agent then sends an RSVP RESV packet (containing the flowspec and other information) to the previous hop.

Each router or host along the path back to the sender receives this RESV packet, uses the flowspec to install the appropriate reservation (if possible), and forwards

the RESV to its previous hop. In this way, each RSVP-capable router or host along the data path installs the reservation according to its capabilities. At the sender, the RSVP Agent passes the RESV packet information to the sender application, which then has information that indicates the actual reservation in place. The sender might choose to wait for the reservation to be in place, or might begin sending data before this happens (although such data is treated by the network as though no reservation were in place). Any router or host that is incapable of supporting the requested reservation might send an error to the receiver, which is then free to perhaps try a lesser reservation.

The z/OS UNIX RSVP agent can provide actual resource reservations on ATM interfaces. The RSVP agent passes the reservation request to the TCP/IP stack, where a bandwidth reserved SVC is established on the ATM link to support the reservation request. The RSVP agent can also cause the Type of Service (TOS) byte to be set for any given RSVP flow, by using the OutgoingTOS keyword on the PolicyAction statement of a defined service policy.

Gathering diagnostic information

The RSVP Agent writes logging information to a log file. The level of logged information is controlled by the LogLevel configuration statement. By default, only error and warning messages are written. To gather more diagnostic information, you can specify a LogLevel value. The maximum information is logged with a LogLevel value of 511. See z/OS Communications Server: IP Configuration Guide for more details on using LogLevel, as well as the location of the log file.

The following information can also be useful in diagnosing RSVP Agent problems:

- Output from the TSO NETSTAT SLAP or **netstat -j** commands
- Output from the **pasearch** command for RSVP scoped policies
- SNMP output from walks of the Network SLAPM2 Subagent MIB tables
- TCP/IP CTRACE output, using the INTERNET and IOCTL CTRACE options
- Policy Agent log output if RSVP scoped policies are defined

Diagnosing RSVP agent problems

Problems with the RSVP agent generally fall into one of the following categories:

- Initialization Problems
- Application Problems
- Service Policy Problems

Initialization problems

Before you begin, if the RSVP Agent does not complete initialization, run it with LogLevel set to 511 and check the log file for error conditions.

Common problems are listed in Table 56:

Table 56. Common RSVP initialization problems

Problem	Cause or action
RSVP Agent not authorized to security product	The RSVP Agent must be authorized to a security product profile. See z/OS Communications Server: IP Configuration Guide for details on setting up the proper authorization.

Table 56. Common RSVP initialization problems (continued)

Problem	Cause or action
Unable to read configuration file	Is the correct configuration file specified? See z/OS Communications Server: IP Configuration Guide for the search order used to locate the configuration file. Does the file exist? Are the permission bits correctly set for a z/OS UNIX file system file?
Unable to associate with the TCP/IP stack	Is the associated TCP/IP stack started? The RSVP Agent uses the TCP/IP image name specified in the configuration file, or uses the standard resolver search order, to locate the name of the TCP/IP stack. The log file indicates the stack name being used.
Unable to initialize interfaces	The RSVP Agent needs to initialize each interface for which it is configured. A pair of "mailboxes" are created for each interface. Check for error messages while creating the "rsvp" and "rsvp-udp" mailboxes for each interface. An error received while trying to join a multicast group on an interface that is not multicast capable is expected, and looks like: `WARNING:.....mailslot_create: setsockopt(MCAST_ADD) failed - EDC5121I Invalid argument.`

Application problems

Before you begin, determine whether a Qos-aware application using RAPI is experiencing problems.

If so, check the items listed in Table 57.

Table 57. RSVP application problems

Problem	Cause or action
RAPI DLL not found	An application using RAPI must have access to the RAPI DLL at run time. This is normally accomplished with the LIBPATH environment variable. Check that the LIBPATH environment variable is specified and that it contains the directory in which the RAPI DLL (rapi.dll) resides, which should be /usr/lib.
Error RAPI_ERR_NORSVP received	If the application receives a RAPI_ERR_NORSVP error code when calling a RAPI function, ensure that the RSVP Agent has been successfully started.

Policy problems

Before you begin, determine whether you are having problems with policies with RSVP scope. Policies with RSVP scope can be defined and made available by way of the Policy Agent.

If problems are encountered using such policies, check the items listed in Table 58.

Table 58. RSVP policy problems

Problem	Cause or action
RSVP policies not being applied to data flows	If the limits imposed by defined RSVP-scoped policies are not taking effect, check that the Policy Agent has been successfully started. The Policy Agent must be active in order for the RSVP Agent to retrieve these policies. Check that the policies are correctly defined. For example, do not specify both inbound and outbound interfaces in a single policy condition because such a policy never maps to any traffic on an end host node. Also, check both the RSVP Agent and Policy Agent log files for errors dealing with obtaining policies.
Policy values not being used or are incorrect	If the values being used in the policies to limit Tspec and Rspec values do not appear to be correct, or do not seem to be applied to RSVP data traffic, be aware that the service policy and Tspec/Rspec units of measure are different. Specifically, the following are different: If the Service Policy Unit is: • MaxRatePerFlow: kilobits/second, the Tspec/Rspec Unit is r/R: bytes/second If the Service Policy Unit is: • MaxTokenBucketPerFlow: kilobits, the Tspec/Rspec Unit is b: bytes To arrive at the values to specify on the service policy, multiply the target Tspec/Rspec value by 8, then divide by 1000. For example, if the target Tspec b value is 6000, the corresponding MaxTokenBucketPerFlow value is 48 (6000 x 8 / 1000 = 48). See Chapter 25, "Diagnosing Policy Agent problems," on page 657 for more information about Policy Agent.

Example log file

Figure 91 on page 691 demonstrates some of the RSVP Agent processing. This log file was created using a LogLevel of 511.

Lines with numbers displayed like **1** are annotations that are described following the log.

```
01
03/22 08:51:01 INFO   :.main: *************** RSVP Agent started ***************
02
03/22 08:51:01 INFO   :...locate_configFile: Specified configuration file: /u/user10/rsvpd1.conf
03/22 08:51:01 INFO   :.main: Using log level 511
03/22 08:51:01 INFO   :..settcpimage: Get TCP images rc - EDC8112I Operation not supported on socket.
03
03/22 08:51:01 INFO   :..settcpimage: Associate with TCP/IP image name = TCPCS
03/22 08:51:02 INFO   :..reg_process: registering process with the system
03/22 08:51:02 INFO   :..reg_process: attempt OS/390 registration
03/22 08:51:02 INFO   :..reg_process: return from registration rc=0
04
03/22 08:51:06 TRACE  :...read_physical_netif: Home list entries returned = 7
03/22 08:51:06 INFO   :...read_physical_netif: index #0, interface VLINK1 has address 129.1.1.1, ifidx 0
03/22 08:51:06 INFO   :...read_physical_netif: index #1, interface TR1 has address 9.37.65.139, ifidx 1
03/22 08:51:06 INFO   :...read_physical_netif: index #2, interface LINK11 has address 9.67.100.1, ifidx 2
03/22 08:51:06 INFO   :...read_physical_netif: index #3, interface LINK12 has address 9.67.101.1, ifidx 3
03/22 08:51:06 INFO   :...read_physical_netif: index #4, interface CTCD0 has address 9.67.116.98, ifidx 4
03/22 08:51:06 INFO   :...read_physical_netif: index #5, interface CTCD2 has address 9.67.117.98, ifidx 5
03/22 08:51:06 INFO   :...read_physical_netif: index #6, interface LOOPBACK has address 127.0.0.1, ifidx 0
03/22 08:51:06 INFO   :....mailslot_create: creating mailslot for timer
03/22 08:51:06 INFO   :...mailbox_register: mailbox allocated for timer
05
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP via UDP
06
03/22 08:51:06 WARNING:.....mailslot_create: setsockopt(MCAST_ADD) failed - EDC8116I Address not available.
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp-udp
03/22 08:51:06 TRACE  :..entity_initialize: interface 129.1.1.1, entity for rsvp allocated and initialized
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP via UDP
03/22 08:51:06 WARNING:.....mailslot_create: setsockopt(MCAST_ADD) failed - EDC8116I Address not available.
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp-udp
03/22 08:51:06 TRACE  :..entity_initialize: interface 9.37.65.139, entity for rsvp allocated and
initialized
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP via UDP
03/22 08:51:06 WARNING:.....mailslot_create: setsockopt(MCAST_ADD) failed - EDC8116I Address not available.
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp-udp
03/22 08:51:06 TRACE  :..entity_initialize: interface 9.67.100.1, entity for rsvp allocated and initialized
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP via UDP
03/22 08:51:06 WARNING:.....mailslot_create: setsockopt(MCAST_ADD) failed - EDC8116I Address not available.
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp-udp
03/22 08:51:06 TRACE  :..entity_initialize: interface 9.67.101.1, entity for rsvp allocated and initialized
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP via UDP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp-udp
03/22 08:51:06 TRACE  :..entity_initialize: interface 9.67.116.98, entity for rsvp allocated and
initialized
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP via UDP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp-udp
03/22 08:51:06 TRACE  :..entity_initialize: interface 9.67.117.98, entity for rsvp allocated and
initialized
```

Figure 91. RSVP Agent processing log

```
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp
03/22 08:51:06 INFO   :.....mailslot_create: creating mailslot for RSVP via UDP
03/22 08:51:06 INFO   :....mailbox_register: mailbox allocated for rsvp-udp
03/22 08:51:06 TRACE  :..entity_initialize: interface 127.0.0.1, entity for rsvp allocated and initialized
03/22 08:51:06 INFO   :......mailslot_create: creating socket for querying route
03/22 08:51:06 INFO   :.....mailbox_register: no mailbox necessary for forward
03/22 08:51:06 INFO   :......mailslot_create: creating mailslot for route engine - informational socket
03/22 08:51:06 TRACE  :......mailslot_create: ready to accept informational socket connection
```

```
03/22 08:51:11 INFO   :.....mailbox_register: mailbox allocated for route
03/22 08:51:11 INFO   :.....mailslot_create: creating socket for traffic control module
03/22 08:51:11 INFO   :....mailbox_register: no mailbox necessary for traffic-control
03/22 08:51:11 INFO   :....mailslot_create: creating mailslot for RSVP client API
03/22 08:51:11 INFO   :...mailbox_register: mailbox allocated for rsvp-api
03/22 08:51:11 INFO   :...mailslot_create: creating mailslot for terminate
03/22 08:51:11 INFO   :..mailbox_register: mailbox allocated for terminate
03/22 08:51:11 INFO   :...mailslot_create: creating mailslot for dump
03/22 08:51:11 INFO   :..mailbox_register: mailbox allocated for dump
03/22 08:51:11 INFO   :...mailslot_create: creating mailslot for (broken) pipe
03/22 08:51:11 INFO   :..mailbox_register: mailbox allocated for pipe
 07
03/22 08:51:11 INFO   :.main: rsvpd initialization complete
 08
03/22 08:52:50 INFO   :......rsvp_api_open: accepted a new connection for rapi
03/22 08:52:50 INFO   :.......mailbox_register: mailbox allocated for mailbox
03/22 08:52:50 TRACE  :......rsvp_event_mapSession: Session=9.67.116.99:1047:6 does not exist
 09
03/22 08:52:50 EVENT  :.....api_reader: api request SESSION
 10
03/22 08:52:50 TRACE  :......rsvp_event_establishSession: local node will send
03/22 08:52:50 INFO   :........router_forward_getOI: Ioctl to get route entry successful
03/22 08:52:50 TRACE  :........router_forward_getOI:         source address:   9.67.116.98
03/22 08:52:50 TRACE  :........router_forward_getOI:         out inf:   9.67.116.98
03/22 08:52:50 TRACE  :........router_forward_getOI:         gateway:   0.0.0.0
03/22 08:52:50 TRACE  :........router_forward_getOI:         route handle:   7f5251c8
 11
03/22 08:52:50 TRACE  :.......event_establishSessionSend: found outgoing if=9.67.116.98 through
forward engine
03/22 08:52:50 TRACE  :......rsvp_event_mapSession: Session=9.67.116.99:1047:6 exists
 12
03/22 08:52:50 EVENT  :.....api_reader: api request SENDER
 13
03/22 08:52:50 INFO   :.......init_policyAPI: papi_debug:  Entering

03/22 08:52:50 INFO   :.......init_policyAPI: papi_debug:  papiLogFunc = 98681F0 papiUserValue = 0

03/22 08:52:50 INFO   :.......init_policyAPI: papi_debug:  Exiting

03/22 08:52:50 INFO   :.......init_policyAPI: APIInitialize:  Entering

03/22 08:52:50 INFO   :.......init_policyAPI: open_socket:  Entering

03/22 08:52:50 INFO   :.......init_policyAPI: open_socket:  Exiting

03/22 08:52:50 INFO   :.......init_policyAPI: APIInitialize:  ApiHandle = 98BDFB0,  connfd = 22

03/22 08:52:50 INFO   :.......init_policyAPI: APIInitialize:  Exiting

03/22 08:52:50 INFO   :.......init_policyAPI: RegisterWithPolicyAPI:  Entering
03/22 08:52:50 INFO   :.......init_policyAPI: RegisterWithPolicyAPI:  Writing to socket = 22

03/22 08:52:50 INFO   :.......init_policyAPI: ReadBuffer:  Entering

03/22 08:52:51 INFO   :.......init_policyAPI: ReadBuffer:  Exiting

03/22 08:52:51 INFO   :.......init_policyAPI: RegisterWithPolicyAPI:  Exiting
03/22 08:52:51 INFO   :.......init_policyAPI: Policy API initialized
03/22 08:52:51 INFO   :......rpapi_getPolicyData: RSVPFindActionName:  Entering

03/22 08:52:51 INFO   :......rpapi_getPolicyData: ReadBuffer:  Entering

03/22 08:52:51 INFO   :......rpapi_getPolicyData: ReadBuffer:  Exiting

03/22 08:52:51 INFO   :......rpapi_getPolicyData: RSVPFindActionName:  Result = 0

03/22 08:52:51 INFO   :......rpapi_getPolicyData: RSVPFindActionName:  Exiting
```

```
14
03/22 08:52:51 INFO   :......rpapi_getPolicyData: found action name CLCat2 for
flow[sess=9.67.116.99:1047:6,source=9.67.116.98:8000]
03/22 08:52:51 INFO   :......rpapi_getPolicyData: RSVPFindServiceDetailsOnActName:  Entering

03/22 08:52:51 INFO   :......rpapi_getPolicyData: ReadBuffer:  Entering

03/22 08:52:51 INFO   :......rpapi_getPolicyData: ReadBuffer:  Exiting

03/22 08:52:51 INFO   :......rpapi_getPolicyData: RSVPFindServiceDetailsOnActName:  Result = 0

03/22 08:52:51 INFO   :......rpapi_getPolicyData: RSVPFindServiceDetailsOnActName:  Exiting

03/22 08:52:51 INFO   :.....api_reader: appl chose service type 1
03/22 08:52:51 INFO   :......rpapi_getSpecData: RSVPGetTSpec:  Entering

03/22 08:52:51 INFO   :......rpapi_getSpecData: RSVPGetTSpec:  Result = 0

03/22 08:52:51 INFO   :......rpapi_getSpecData: RSVPGetTSpec:  Exiting

03/22 08:52:51 TRACE  :.....api_reader: new service=1, old service=0
03/22 08:52:51 INFO   :.......rsvp_flow_stateMachine: state SESSIONED, event PATHDELTA
15
03/22 08:52:51 TRACE  :........rsvp_action_nHop: constructing a PATH
03/22 08:52:51 TRACE  :........flow_timer_start: started T1
16
03/22 08:52:51 TRACE  :.......rsvp_flow_stateMachine: entering state PATHED
03/22 08:52:51 TRACE  :........mailslot_send: sending to (9.67.116.99:0)
03/22 08:52:51 TRACE  :........mailslot_send: sending to (9.67.116.99:1698)
17
03/22 08:52:51 TRACE  :.....rsvp_event: received event from RAW-IP on interface 9.67.116.98
03/22 08:52:51 TRACE  :......rsvp_explode_packet: v=1,flg=0,type=2,cksm=54875,ttl=255,rsv=0,len=84
03/22 08:52:51 TRACE  :.......rsvp_parse_objects: STYLE is WF
03/22 08:52:51 INFO   :.......rsvp_parse_objects: obj RSVP_HOP hop=9.67.116.99, lih=0
03/22 08:52:51 TRACE  :......rsvp_event_mapSession: Session=9.67.116.99:1047:6 exists
03/22 08:52:51 INFO   :.......rsvp_flow_stateMachine: state PATHED, event RESVDELTA
18
03/22 08:52:51 TRACE  :........traffic_action_oif: is to install filter
03/22 08:52:51 INFO   :.........qosmgr_request: src-9.67.116.98:8000 dst-9.67.116.99:1047 proto-6
rthdl-7f5251c8
19
03/22 08:52:51 INFO   :.........qosmgr_request: [CL r=90000 b=6000 p=110000 m=1024 M=2048]
03/22 08:52:51 INFO   :.........qosmgr_request: Ioctl to add reservation successful
03/22 08:52:51 INFO   :..........rpapi_Reg_UnregFlow: RSVPPutActionName:  Entering
03/22 08:52:51 INFO   :..........rpapi_Reg_UnregFlow: ReadBuffer:  Entering

03/22 08:52:52 INFO   :..........rpapi_Reg_UnregFlow: ReadBuffer:  Exiting

03/22 08:52:52 INFO   :..........rpapi_Reg_UnregFlow: RSVPPutActionName:  Result = 0

03/22 08:52:52 INFO   :..........rpapi_Reg_UnregFlow: RSVPPutActionName:  Exiting

03/22 08:52:52 INFO   :..........rpapi_Reg_UnregFlow: flow[sess=9.67.116.99:1047:6,
source=9.67.116.98:8000] registered with CLCat2
03/22 08:52:52 EVENT  :..........qosmgr_response: RESVRESP from qosmgr, reason=0, qoshandle=8b671d0
03/22 08:52:52 INFO   :..........qosmgr_response: src-9.67.116.98:8000 dst-9.67.116.99:1047 proto-6
03/22 08:52:52 TRACE  :...........traffic_reader: tc response msg=1, status=1
03/22 08:52:52 INFO   :...........traffic_reader: Reservation req successful[session=9.67.116.99:1047:6,
source=9.67.116.98:8000, qoshd=8b671d0]
20
03/22 08:52:52 TRACE  :........api_action_sender: constructing a RESV
03/22 08:52:52 TRACE  :........flow_timer_stop: stopped T1
03/22 08:52:52 TRACE  :........flow_timer_stop: Stop T4
03/22 08:52:52 TRACE  :........flow_timer_start: started T1
03/22 08:52:52 TRACE  :........flow_timer_start: Start T4
21
```

```
03/22 08:52:52 TRACE  :.......rsvp_flow_stateMachine: entering state RESVED
 22
03/22 08:53:07 EVENT  :..mailslot_sitter: process received signal SIGALRM
03/22 08:53:07 TRACE  :.....event_timerT1_expire: T1 expired
03/22 08:53:07 INFO   :......router_forward_getOI: Ioctl to query route entry successful
03/22 08:53:07 TRACE  :......router_forward_getOI:         source address:   9.67.116.98
03/22 08:53:07 TRACE  :......router_forward_getOI:         out inf:   9.67.116.98
03/22 08:53:07 TRACE  :......router_forward_getOI:         gateway:   0.0.0.0
03/22 08:53:07 TRACE  :......router_forward_getOI:         route handle:   7f5251c8
03/22 08:53:07 INFO   :......rsvp_flow_stateMachine: state RESVED, event T1OUT
03/22 08:53:07 TRACE  :.......rsvp_action_nHop: constructing a PATH
03/22 08:53:07 TRACE  :.......flow_timer_start: started T1
03/22 08:53:07 TRACE  :......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:53:07 TRACE  :.......mailslot_send: sending to (9.67.116.99:0)
 23
03/22 08:53:22 TRACE  :.....rsvp_event: received event from RAW-IP on interface 9.67.116.98
03/22 08:53:22 TRACE  :......rsvp_explode_packet: v=1,flg=0,type=2,cksm=54875,ttl=255,rsv=0,len=84
03/22 08:53:22 TRACE  :.......rsvp_parse_objects: STYLE is WF
03/22 08:53:22 INFO   :.......rsvp_parse_objects: obj RSVP_HOP hop=9.67.116.99, lih=0
03/22 08:53:22 TRACE  :......rsvp_event_mapSession: Session=9.67.116.99:1047:6 exists
03/22 08:53:22 INFO   :.......rsvp_flow_stateMachine: state RESVED, event RESV
03/22 08:53:22 TRACE  :........flow_timer_stop: Stop T4
03/22 08:53:22 TRACE  :........flow_timer_start: Start T4
03/22 08:53:22 TRACE  :.......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:53:22 EVENT  :..mailslot_sitter: process received signal SIGALRM
03/22 08:53:22 TRACE  :.....event_timerT1_expire: T1 expired
03/22 08:53:22 INFO   :......router_forward_getOI: Ioctl to query route entry successful
03/22 08:53:22 TRACE  :......router_forward_getOI:         source address:   9.67.116.98
03/22 08:53:22 TRACE  :......router_forward_getOI:         out inf:   9.67.116.98
03/22 08:53:22 TRACE  :......router_forward_getOI:         gateway:   0.0.0.0
03/22 08:53:22 TRACE  :......router_forward_getOI:         route handle:   7f5251c8
03/22 08:53:22 INFO   :......rsvp_flow_stateMachine: state RESVED, event T1OUT
03/22 08:53:22 TRACE  :.......rsvp_action_nHop: constructing a PATH
03/22 08:53:22 TRACE  :.......flow_timer_start: started T1
03/22 08:53:22 TRACE  :......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:53:22 TRACE  :.......mailslot_send: sending to (9.67.116.99:0)
03/22 08:53:38 EVENT  :..mailslot_sitter: process received signal SIGALRM
03/22 08:53:38 TRACE  :.....event_timerT1_expire: T1 expired
03/22 08:53:38 INFO   :......router_forward_getOI: Ioctl to query route entry successful
03/22 08:53:38 TRACE  :......router_forward_getOI:         source address:   9.67.116.98
03/22 08:53:38 TRACE  :......router_forward_getOI:         out inf:   9.67.116.98
03/22 08:53:38 TRACE  :......router_forward_getOI:         gateway:   0.0.0.0
03/22 08:53:38 TRACE  :......router_forward_getOI:         route handle:   7f5251c8
03/22 08:53:38 INFO   :......rsvp_flow_stateMachine: state RESVED, event T1OUT
03/22 08:53:38 TRACE  :.......rsvp_action_nHop: constructing a PATH
03/22 08:53:38 TRACE  :.......flow_timer_start: started T1
03/22 08:53:38 TRACE  :......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:53:38 TRACE  :.......mailslot_send: sending to (9.67.116.99:0)
03/22 08:53:52 TRACE  :.....rsvp_event: received event from RAW-IP on interface 9.67.116.98
03/22 08:53:52 TRACE  :......rsvp_explode_packet: v=1,flg=0,type=2,cksm=54875,ttl=255,rsv=0,len=84
03/22 08:53:52 TRACE  :.......rsvp_parse_objects: STYLE is WF
03/22 08:53:52 INFO   :.......rsvp_parse_objects: obj RSVP_HOP hop=9.67.116.99, lih=0
03/22 08:53:52 TRACE  :......rsvp_event_mapSession: Session=9.67.116.99:1047:6 exists
03/22 08:53:52 INFO   :.......rsvp_flow_stateMachine: state RESVED, event RESV
03/22 08:53:52 TRACE  :........flow_timer_stop: Stop T4
03/22 08:53:52 TRACE  :........flow_timer_start: Start T4
03/22 08:53:52 TRACE  :.......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:53:53 EVENT  :..mailslot_sitter: process received signal SIGALRM
03/22 08:53:53 TRACE  :.....event_timerT1_expire: T1 expired
03/22 08:53:53 INFO   :......router_forward_getOI: Ioctl to query route entry successful
03/22 08:53:53 TRACE  :......router_forward_getOI:         source address:   9.67.116.98
03/22 08:53:53 TRACE  :......router_forward_getOI:         out inf:   9.67.116.98
03/22 08:53:53 TRACE  :......router_forward_getOI:         gateway:   0.0.0.0
03/22 08:53:53 TRACE  :......router_forward_getOI:         route handle:   7f5251c8
03/22 08:53:53 INFO   :......rsvp_flow_stateMachine: state RESVED, event T1OUT
03/22 08:53:53 TRACE  :.......rsvp_action_nHop: constructing a PATH
03/22 08:53:53 TRACE  :.......flow_timer_start: started T1
```

```
03/22 08:53:53 TRACE  :......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:53:53 TRACE  :.......mailslot_send: sending to (9.67.116.99:0)
03/22 08:54:09 EVENT  :..mailslot_sitter: process received signal SIGALRM
03/22 08:54:09 TRACE  :.....event_timerT1_expire: T1 expired
03/22 08:54:09 INFO   :......router_forward_getOI: Ioctl to query route entry successful
03/22 08:54:09 TRACE  :......router_forward_getOI:         source address:   9.67.116.98
03/22 08:54:09 TRACE  :......router_forward_getOI:         out inf:   9.67.116.98
03/22 08:54:09 TRACE  :......router_forward_getOI:         gateway:   0.0.0.0
03/22 08:54:09 TRACE  :......router_forward_getOI:         route handle:   7f5251c8
03/22 08:54:09 INFO   :......rsvp_flow_stateMachine: state RESVED, event T1OUT
03/22 08:54:09 TRACE  :.......rsvp_action_nHop: constructing a PATH
03/22 08:54:09 TRACE  :.......flow_timer_start: started T1
03/22 08:54:09 TRACE  :......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:54:09 TRACE  :.......mailslot_send: sending to (9.67.116.99:0)
03/22 08:54:22 TRACE  :.....rsvp_event: received event from RAW-IP on interface 9.67.116.98
03/22 08:54:22 TRACE  :......rsvp_explode_packet: v=1,flg=0,type=2,cksm=54875,ttl=255,rsv=0,len=84
03/22 08:54:22 TRACE  :.......rsvp_parse_objects: STYLE is WF
03/22 08:54:22 INFO   :.......rsvp_parse_objects: obj RSVP_HOP hop=9.67.116.99, lih=0
03/22 08:54:22 TRACE  :......rsvp_event_mapSession: Session=9.67.116.99:1047:6 exists
03/22 08:54:22 INFO   :.......rsvp_flow_stateMachine: state RESVED, event RESV
03/22 08:54:22 TRACE  :........flow_timer_stop: Stop T4
03/22 08:54:22 TRACE  :........flow_timer_start: Start T4
03/22 08:54:22 TRACE  :.......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:54:24 EVENT  :..mailslot_sitter: process received signal SIGALRM
03/22 08:54:24 TRACE  :.....event_timerT1_expire: T1 expired
03/22 08:54:24 INFO   :......router_forward_getOI: Ioctl to query route entry successful
03/22 08:54:24 TRACE  :......router_forward_getOI:         source address:   9.67.116.98
03/22 08:54:24 TRACE  :......router_forward_getOI:         out inf:   9.67.116.98
03/22 08:54:24 TRACE  :......router_forward_getOI:         gateway:   0.0.0.0
03/22 08:54:24 TRACE  :......router_forward_getOI:         route handle:   7f5251c8
03/22 08:54:24 INFO   :......rsvp_flow_stateMachine: state RESVED, event T1OUT
03/22 08:54:24 TRACE  :.......rsvp_action_nHop: constructing a PATH
03/22 08:54:24 TRACE  :.......flow_timer_start: started T1
03/22 08:54:24 TRACE  :......rsvp_flow_stateMachine: reentering state RESVED
03/22 08:54:24 TRACE  :.......mailslot_send: sending to (9.67.116.99:0)
03/22 08:54:35 TRACE  :......rsvp_event_mapSession: Session=9.67.116.99:1047:6 exists
 24
03/22 08:54:35 EVENT  :.....api_reader: api request SENDER_WITHDRAW
03/22 08:54:35 INFO   :.......rsvp_flow_stateMachine: state RESVED, event PATHTEAR
 25
03/22 08:54:35 TRACE  :........traffic_action_oif: is to remove filter
03/22 08:54:35 INFO   :.........qosmgr_request: Ioctl to remove reservation successful
03/22 08:54:35 INFO   :..........rpapi_Reg_UnregFlow: RSVPRemActionName:  Entering

03/22 08:54:35 INFO   :..........rpapi_Reg_UnregFlow: ReadBuffer:  Entering

03/22 08:54:35 INFO   :..........rpapi_Reg_UnregFlow: ReadBuffer:  Exiting

03/22 08:54:35 INFO   :..........rpapi_Reg_UnregFlow: RSVPRemActionName:  Result = 0

03/22 08:54:35 INFO   :..........rpapi_Reg_UnregFlow: RSVPRemActionName:  Exiting

03/22 08:54:35 INFO   :..........rpapi_Reg_UnregFlow: flow[sess=9.67.116.99:1047:6,
source=9.67.116.98:8000] unregistered from CLCat2
03/22 08:54:35 EVENT  :..........qosmgr_response: DELRESP from qosmgr, reason=0, qoshandle=0
03/22 08:54:35 INFO   :..........qosmgr_response: src-9.67.116.98:8000 dst-9.67.116.99:1047 proto-6
03/22 08:54:35 TRACE  :...........traffic_reader: tc response msg=3, status=1
 26
03/22 08:54:35 TRACE  :........rsvp_action_nHop: constructing a PATHTEAR
03/22 08:54:35 TRACE  :........flow_timer_stop: stopped T1
03/22 08:54:35 TRACE  :........flow_timer_stop: Stop T4
 27
03/22 08:54:35 TRACE  :.......rsvp_flow_stateMachine: entering state SESSIONED
03/22 08:54:35 TRACE  :........mailslot_send: sending to (9.67.116.99:0)
03/22 08:54:35 TRACE  :......rsvp_event_propagate: flow[session=9.67.116.99:1047:6,
source=9.67.116.98:8000] ceased
 28
```

```
03/22 08:54:35 EVENT  :.....api_reader: api request CLOSE
03/22 08:54:35 INFO   :.......rsvp_flow_stateMachine: state SESSIONED, event PATHTEAR
03/22 08:54:35 PROTERR:.......rsvp_flow_stateMachine: state SESSIONED does not expect event PATHTEAR
 29
03/22 08:54:53 EVENT  :..mailslot_sitter: process received signal SIGTERM
03/22 08:54:53 INFO   :...check_signals: received TERM signal
03/22 08:54:53 INFO   :......term_policyAPI: UnRegisterFromPolicyAPI:  Entering

03/22 08:54:53 INFO   :......term_policyAPI: ReadBuffer:  Entering

03/22 08:54:53 INFO   :......term_policyAPI: ReadBuffer:  Exiting

03/22 08:54:53 INFO   :......term_policyAPI: UnRegisterFromPolicyAPI:  Result = 0

03/22 08:54:53 INFO   :......term_policyAPI: UnRegisterFromPolicyAPI:  Exiting

03/22 08:54:53 INFO   :......term_policyAPI: APITerminate:  Entering

03/22 08:54:53 INFO   :......term_policyAPI: APITerminate:  Exiting

03/22 08:54:53 INFO   :......term_policyAPI: Policy API terminated
03/22 08:54:53 INFO   :......dreg_process: deregistering process with the system
03/22 08:54:53 INFO   :......dreg_process: attempt to dereg (ifaeddrg_byaddr)
03/22 08:54:53 INFO   :......dreg_process: rc from ifaeddrg_byaddr  rc =0
03/22 08:54:53 INFO   :.....terminator: process terminated with exit code 0
```

Following are short descriptions of the numbered items in the trace:

01 The RSVP Agent is started.

02 The configuration file being used is reported.

03 The name of the TCP/IP stack that the RSVP Agent associates itself with is reported.

04 The name and IP address of the interfaces configured to the associated stack are reported. Note that the RSVP Agent gets notified by the stack of any interface additions, deletions, or changes after this point.

05 The interfaces are initialized one by one.

06 Some interface types are not enabled for multicasting. Therefore, when the RSVP Agent tries to enable multicasting, a warning is reported. Such interfaces can still be used for unicasting.

07 RSVP Agent initialization is complete.

08 An application makes its first RAPI call, initializing the RAPI interface with the RSVP Agent.

09 The type of RAPI request is SESSION, meaning a rapi_session() call was made.

10 The RSVP Agent determines what the application sends based on the specified destination address not being a local interface.

11 The outbound interface to use for the session is returned from the stack.

12 The application issues a rapi_sender() call, passing the Tspec.

13 The Policy Agent interface is initialized.

14 The policy action "CLCat2" is obtained from the Policy Agent for the specified flow.

15 The RSVP Agent constructs an RSVP PATH packet to be sent to the destination.

16 The flow enters the pathed stated (PATHED), meaning a PATH packet has been sent for the flow.

17 An RSVP RESV packet is received from the RSVP Agent at the receiver node, specifying the reservation parameters.

18 The RSVP Agent installs the reservation request into the TCP/IP stack and registers the flow with the Policy Agent.

19 The type of reservation request is shown (CL, for Controlled Load) along with the reservation parameters (the r, b, p, m, M values in Tspec format).

20 The RESV packet values are passed to the sender application.

21 The flow enters the reserved state (RESVED), meaning the reservation has been put in place and the RESV packet has been forwarded to the previous hop (in this case the sender application).

22 A T1 timeout occurs, meaning a PATH refresh packet is sent. This occurs every 15 seconds.

23 A refreshed RESV packet is received from the RSVP Agent at the receiver node. This occurs every 30 seconds.

24 The application issues a rapi_release() call to end the RAPI session.

25 The reservation is removed from the TCP/IP stack and unregistered from the Policy Agent.

26 A PATHTEAR packet is constructed and sent, to tear down the flow along the data path.

27 The flow enters the sessioned state (SESSIONED), meaning that the flow has been torn down.

28 The application closes the API session, resulting in an error being reported because the state of the flow is SESSIONED. This error can be ignored.

29 A SIGTERM signal is received (due to a **kill** command issued from the UNIX shell), and the RSVP Agent shuts itself down.

Chapter 27. Diagnosing intrusion detection problems

This topic provides information and guidance to diagnose Intrusion Detection Service (IDS) problems, including traffic regulation management daemon (TRMD) related problems. It contains the following subtopics:

- "Overview"
- "Diagnosing IDS policy problems"
- "Diagnosing IDS output problems" on page 700
- "Diagnosing TRMD problems" on page 704
- "Documentation for the IBM Software Support Center" on page 704

Overview

The Intrusion Detection Services policy is installed into the stack by the Policy Agent (PAGENT). After the policy is installed, IDS detects, processes, and reports on events as requested by the policy. TRMD, part of IDS, handles reporting IDS statistics and events to syslogd. Problems might occur in the following areas:

- Policy installation
- Output to syslogd, the console, or the IDS trace missing or volume too high
- TRMD initialization

Diagnosing IDS policy problems

This topic describes the commands used to diagnose IDS policy problems.

Some intrusion detection services (IDS) policies are not mapped until they are needed. Attack policies, scan global policies, and scan event policies for protocols ICMP and ICMPv6 are mapped immediately when the policy is installed in the stack. Scan Event policies for protocols TCP and UDP are mapped on the first occurrence of a potentially countable event. TR policies for protocol TCP are mapped when a local application does a listen() and when a client completes the three-way connection handshake. TR policies for protocol UDP are mapped when an inbound datagram arrives for a bound port.

Step for determining which IDS policies are active in Policy Agent

Determine which IDS policies are active in Policy Agent by using the **pasearch -i** command.

Before you begin

If you are running multiple stacks, ensure that **pasearch** is reporting on the stack you are interested in.

Procedure

1. Use **pasearch -i** (see z/OS Communications Server: IP System Administrator's Commands) to see what IDS policies are active in Policy Agent.
2. See Chapter 25, "Diagnosing Policy Agent problems," on page 657 if you do not see the IDS policies expected.

© Copyright IBM Corp. 2000, 2015

Step for determining how your IDS policies have been mapped by the stack

Determine how your IDS policies were mapped by the stack by using the Netstat command.

Before you begin

If you are running multiple stacks, ensure that your resolver configuration correctly identifies the stack you are interested in. Ensure that your IDS policies are correctly defined.

Procedure

- Use NETSTAT IDS or **netstat -k** (see z/OS Communications Server: IP System Administrator's Commands) to see how your IDS policies have been mapped by the stack.
- See IDS policy considerations in z/OS Communications Server: IP Configuration Guide.

Diagnosing IDS output problems

The following topics describe diagnostic steps for some problems you might encounter.

Steps for determining why IDS syslogd output is missing

Determine the cause for missing IDS syslogd output.

Procedure

Perform the following steps:

1. Ensure that Policy Agent is running on this system
2. Ensure that TRMD is running for this stack on this system. Consider using TCPIP PROFILE Autolog for TRMD.
3. Ensure that syslogd is running on this system.
4. Ensure that syslogd is configured for IDS output:
   - TRMD always writes to the syslog daemon facility.
   - Events are written to the syslog level configured in the relevant policy. Statistics are always written to INFO level.
   - If running multiple TRMDs, consider using trmd jobname prefix to separate IDS output by stack.

IDS console output

Under certain conditions, IDS suppresses console messages to avoid flooding the system console.

Scan detection is reported at most once per fast scan interval for a particular source IP address. If a scan is continually detected for the same source IP address, consider adding this address to your scan exclusion list (if this user is legitimately accessing resources). The installation also has the option of requesting notification to syslogd rather than to the console. The same criteria is used for reporting scans to syslogd as to the console.

IDS attack policy actions support the maximum event message parameter. If specified, this limits the number of times the same attack type is reported to the system console within any 5-minute time period.

Traffic regulation for protocol TCP suppresses console reporting of the following three events that could occur repeatedly.

- Only the first connection denied, when an application exceeds the TR TCP total connections limit, is reported during each port constrained period.
- Only the first connection denied, when a source host exceeds the TR TCP percentage available limit, is reported until the number of connections by that source host to this application drops below 88% of the limit and at least 2 connections below the limit.
- Connections that would exceed the TR TCP percentage of available connections per source host, but are allowed because of a higher value in QoS policy, are reported to syslogd only.

IDS packet trace output

Use the following references or guidelines for IDS packet trace output:

- See "Intrusion Detection Services trace (SYSTCPIS)" on page 149 if message `EZZ4210I CTRACE DEFINE FAILED FOR CTIIDS00` is issued at stack initialization.
- Consider starting the MVS external writer. See "Formatting packet traces using IPCS" on page 96 for information about formatting the IDS packet trace in a dump.
- For IDS attack policy, packets associated with attack events can be traced. For most attack types, a single packet triggers an event and the packet is traced. To prevent trace flooding, a maximum of 100 attack packets per attack type are traced within a 5-minute interval. For the flood attack type, the first 100 packets that are discarded during the flood are traced. For the TCP queue size, global TCP stall, and EE XID flood attack types, no IDS tracing is done.

Unusual conditions

Most messages that are issued by IDS relate to the detection of an IDS condition. However, the messages that are mentioned below should be investigated because they signal conditions which affect IDS normal processing that might result in IDS information being lost or delayed.

Buffer overflow transferring message data between the stack and TRMD

The following messages in syslog indicate that IDS events or statistics are being generated at a rate that is overflowing internal buffers used to relay the messages from the stack to TRMD. These messages are a warning that actual event or statistics messages are missing from the syslog. If these messages occur frequently, then IDS policy changes are necessary to reduce the amount of IDS logging, or the amount of statistics information, being generated.

```
EZZ9325I TRMD Log records missing: logtype,logmissing
EZZ9326I TRMD Statistics records missing: stattype,statmissing
```

A large number of syslogd messages can be generated for the global TCP stall attack type if the syslogd detailed parameter is specified in the configuration. A message (either EZZ8673I or EZZ8674I) is generated for each stalled connection when a global TCP stall condition is detected and syslogd detail information has been requested. Consider updating your policy to turn off the syslogd detailed parameter for the global TCP stall attack type if these messages are causing a buffer overflow.

Repeated attacks of the same type at a high rate

A message is issued in syslog to indicate that attack policy is in place and the attack type indicated is occurring repeatedly at a high rate. To avoid flooding syslog and conserve system resources, a maximum of 100 event messages per attack type are logged to syslogd within a 5-minute interval. This limit is in effect for all attack types except the global TCP stall attack type. The following message indicates the number of duplicate attacks for which messages have been suppressed.

```
EZZ9327I TRMD Attack log records suppressed: attack_type,count
```

Scan storage constrained

The following is an example of a console message issued if scan detection attempted to obtain storage in order to track a potential scan event and could not obtain the required amount of storage.

```
EZZ8761I IDS EVENT DETECTED
EZZ8730I STACK TCPCS
EZZ8762I EVENT TYPE: SCAN STORAGE CONSTRAINED
EZZ8763I CORRELATOR 0 - PROBEID 0300FFF3
EZZ8766I IDS RULE N/A
EZZ8767I IDS ACTION N/A
```

Processing continues without adding the tracking information for this packet or for subsequent packets in the current internal interval (an internal interval is either 30 or 60 seconds). This could result in missing potential scan events.

The installation should attempt to determine the cause of the storage shortage. Scan detection itself can potentially consume large amounts of storage and should be looked at as part of the problem determination. The following are two ways to determine whether scan is consuming large amounts of storage.

- Console message EZZ8768I (`EZZ8768I IDS SCAN STORAGE EXCEEDED nbrmeg MB, TRACKING nbrsip SOURCE IP ADDRESSES`) is issued after scan detection acquires more than a megabyte of storage. This message is reissued at each power of 2 MB increments (for example, 1 MB, 2 MB , 4 MB, 8 MB, and so on).
- The Netstat IDS command displays high-level scan information. For example:

  ```
  SCAN DETECTION:
    GLOBRULENAME: IDS-RULE4
    ICMPRULENAME: IDS-RULE8
    TOTDETECTED:  1           DETCURRPLC: 1
    DETCURRINT:   0           INTERVAL:   30
    SRCIPSTRKD:   125         STRGLEV:    00000M
  ```

  The SRCIPSTRKD field indicates the number of source IPs being tracked and the STRGLEV field indicates the number of megabytes of storage that scan is holding.

If scan processing is contributing to the storage shortage, consider changing the scan policy. If the installation has set the scan sensitivity to HIGH on high usage ports, consider reducing the sensitivity level or removing the port from scan detection until the storage constraint is resolved.

When scan starts to successfully obtain storage again, a `SCAN STORAGE UNCONSTRAINED` message is issued.

Excessive processing time for scans

The following is an example of a console message issued as a result of excessive processing time for scans:

```
EZZ8761I IDS EVENT DETECTED
EZZ8730I STACK TCPCS
EZZ8762I EVENT TYPE: SCAN INTERVAL OVERRUN
EZZ8763I CORRELATOR 0 - PROBEID 0300FFF5
EZZ8766I IDS RULE N/A
EZZ8767I IDS ACTION N/A
```

If an installation repeatedly receives this message, scan processing is not able to complete its evaluation of the source IP addresses it is tracking in its normal interval (either 30 or 60 seconds). This could delay the detection of subsequent scans. This most likely indicates that a large number of source IP addresses are being monitored. If the policy is using high scan sensitivity, the installation should consider lowering the scan sensitivity level for high usage ports.

Interface flood detection disabled

To track data for interface flood detection, private storage is obtained when IDS starts monitoring an interface. If the storage cannot be obtained, IDS is not able to detect an interface flood for the interface. A console message and a syslogd message are issued to report the condition.

The following example shows the console message that is issued:

```
EZZ8761I IDS EVENT DETECTED
EZZ8730I STACK TCPCS
EZZ8762I EVENT TYPE: INTERFACE FLOOD DETECTION DISABLED
EZZ8763I CORRELATOR 20 - PROBEID 04070015
EZZ8770I INTERFACE OSAQDIO4L
EZZ8765I DESTINATION IP ADDRESS 5.72.107.78 - PORT 0
EZZ8766I IDS RULE AttackFlood-rule
EZZ8767I IDS ACTION AttackLog-action
```

The following is an example of the syslogd message:

```
EZZ8658I TRMD ATTACK Interface Flood Detection Disabled:12/23/2002 20:39:35.00,
ifcname=OSAQDIO4L, dipaddr=5.72.107.78,correlator=20,probeid=04070015,
sensorhostname=MVS34.tcp.com
```

These messages indicate a storage constraint has prevented the initialization of interface flood detection for the interface specified in the message. Interface flood detection for other interfaces is not affected.

When the problem causing the storage constraint is resolved, the Interface Flood detection support can be activated by removing the IDS ATTACK FLOOD policy and then adding the IDS ATTACK FLOOD policy again, or by stopping and restarting the interface.

Interface flood storage constrained

The following message in syslogd indicates that private storage needed to collect informational data related to a possible interface flood condition could not be obtained:

```
EZZ8659I TRMD ATTACK Interface Flood storage constrained:timestamp,ifcname=ifcname,
dipaddr=dipaddr,correlator=correlator,probeid=04070016,sensorhostname=sensorhostname
```

The informational data provided by the EZZ8655I and EZZ8656I syslogd messages issued for the interface in the same time period might be incomplete. Collection of informational data for the interface that requires additional storage is temporarily suspended and resumes at the start of the next one-minute interval.

Diagnosing TRMD problems

The most common type of TRMD problem is initialization.

The TRMD writes logging information to a log file. The level of logged information is controlled by the **-d** startup option. To gather more diagnostic information, you can start the TRMD with the **-d** startup option. The maximum information is logged with the **-d 3** option. Log output is directed to the syslog daemon (syslogd). See the z/OS Communications Server: IP Configuration Reference for more details on using the **-d** startup option.

Problems with initialization of the TRMD include:

- Starting TRMD from the console.

  TRMD might fail with an ABEND=S000 U4093 REASON=00000090 because an OMVS segment was not defined for the TRMD ID.

  Check the job output.

```
IEF403I TRMD - STARTED - TIME=12.48.55
ICH408I JOB(TRMD    ) STEP(TRMD    ) CL(PROCESS )
  OMVS SEGMENT NOT DEFINED
IEA995I SYMPTOM DUMP OUTPUT
  USER COMPLETION CODE=4093 REASON CODE=00000090
 TIME=12.48.58  SEQ=00065  CPU=0000  ASID=002B
 PSW AT TIME OF ERROR  078D1000   8000AA7A  ILC 2  INTC 0D
   ACTIVE LOAD MODULE           ADDRESS=00007E70  OFFSET=0000
   NAME=CEEBINIT
   DATA AT PSW  0000AA74 - 00181610  0A0D47F0  B10A1811
   GR 0: 84000000   1: 84000FFD
      2: 00000090   3: 00000001
        4: 0001C2A0   5: 0001C144
        6: 00016560   7: 000169D0
        8: 00000016   9: 098E374E
        A: 00000004   B: 8000A9A8
        C: 00017AC0   D: 0001C018
        E: 00000000   F: 00000090
   END OF SYMPTOM DUMP
  IEF450I TRMD TRMD - ABEND=S000 U4093 REASON=00000090
           TIME=12.48.58
  IEF404I TRMD - ENDED - TIME=12.48.58
  $HASP395 TRMD     ENDED
CEE5101C During initialization, the z/OS Unix System
Services callable service BPX1MSS failed. The system return
code was 0000000156 , the reason code was 0B0C00F9 .
The application will be terminated.
```

  Verify that an OMVS segment exists for TRMD by issuing the TSO **lu** command from a user ID that has authority to issue the LU trmd OMVS command. If an OMVS segment does not exist, use the ALU command to update the user's OMVS data. For example, `ALTUSER trmd  OMVS(UID(0000) HOME('/') PROGRAM('/bin/sh')`.

- The TCP/IP stack is not up and message EZZ8498I is received.

  Verify that the TCP/IP stack is up.

Documentation for the IBM Software Support Center

When contacting the IBM Software Support Center for problem resolution, some or all of the following information might be required:

- Gather TRMD debugging data by starting TRMD with the **trmd -d 3** command. See "Diagnosing TRMD problems."

- Start CTRACE in the stack to gather related information. See “Component trace” on page 45.
- The output from the **pasearch -i** command. See z/OS Communications Server: IP System Administrator's Commands.
- The output from the Netstat IDS/-k command. See z/OS Communications Server: IP System Administrator's Commands.

Chapter 28. Diagnosing Application Transparent Transport Layer Security (AT-TLS)

AT-TLS transparently performs Transport Layer Security (TLS) on behalf of the application by invoking the z/OS System Secure Socket Layer (SSL) in the TCP transport layer. System SSL provides support for the TLS and SSL protocols. AT-TLS uses a policy-based configuration, and the Policy Agent application is required to define rules and actions to the TCP/IP stack for TCP connections using AT-TLS. Displays for AT-TLS policy are provided by **pasearch** and Netstat.

This topic describes how to diagnose AT-TLS problems and includes the following sections:

- "Common AT-TLS startup errors"
- "Steps for diagnosing AT-TLS problems"
- "AT-TLS traces" on page 709
- "AT-TLS return codes" on page 712
- "SIOCTTLSCTL ioctl return codes" on page 718

Common AT-TLS startup errors

The following list describes startup errors, possible causes, and actions to take.

- If message EZZ4248E is written to the console and not released, one of the following might have occurred:
  - Policy Agent has not been started.
  - Policy Agent configuration does not contain a TCPImage statement for this stack, or the stack policy configuration does not contain any local or remote AT-TLS policies.
  - Policy Agent is not permitted to create a socket with this stack. Ensure that the SERVAUTH class is active. Ensure that the EZB.INITSTACK.mvsname.tcpname resource profile is defined and that Policy Agent is permitted to it. If the EZB.STACKACCESS.mvsname.tcpname resource profile is defined, ensure that Policy Agent is permitted to it.
- If message EZD1288I is displayed, the AT-TLS group is configured with FIPS140 on, but ICSF is not active. Start ICSF before starting AT-TLS groups configured with FIPS140 On.
- If applications started after the stack fail to create a socket (errno EAGAIN, errno2 JrTcpNotActive), the stack is probably being configured for AT-TLS, and the application has been started before AT-TLS policy has been installed. If this is a required network infrastructure application, permit it to the EZB.INITSTACK.mvsname.tcpname resource profile in the SERVAUTH class. If it is not a required network infrastructure application, either start it after message EZZ4248E is released or modify the application to wait a short time and try again when the errno is EAGAIN.
- If message `EZD1287I TTLS Error RC: 5020 Group Init` is displayed, the TCP/IP stack was not able to load the System SSL DLL required for AT-TLS processing.

Steps for diagnosing AT-TLS problems

Diagnose AT-TLS problems.

© Copyright IBM Corp. 2000, 2015

Procedure

Perform the following steps:

1. Issue **pasearch -t** to see all AT-TLS policies that are active in Policy Agent. See z/OS Communications Server: IP System Administrator's Commands for more information about the **pasearch -t** command. If you are running multiple stacks, ensure that **pasearch** is reporting on the stack you are interested in. If you do not see the AT-TLS policies that you expected, see z/OS Communications Server: IP System Administrator's Commands for more information about displaying policy based networking information.
2. Issue Netstat TTLS COnn *connid* or Netstat -x COnn *connid* to determine whether the stack mapped a connection to AT-TLS policy and, if so, to which policy it was mapped. For more information about the netstat commands, see z/OS Communications Server: IP System Administrator's Commands. Ensure that your AT-TLS policies are correctly defined. See the AT-TLS information in z/OS Communications Server: IP Configuration Guide and the AT-TLS Policy statements in z/OS Communications Server: IP Configuration Reference for more information about configuring AT-TLS policies.
3. In cases where AT-TLS connections do not map to any policy, verify that TCPCONFIG TTLS has been specified. Netstat configuration shows the current setting of AT-TLS.

   AT-TLS connection mapping is performed based on the following attributes:

   - Local IP Address
   - Remote IP address
   - Local Port
   - Remote Port
   - Direction
   - Job name
   - User ID

   The AT-TLS policy rules are searched, starting with the highest priority rules, for the first match.

   Then the internal SecondaryMap table is searched by process ID and the two IP addresses used on the connection. The SecondaryMap table contains entries for active connections that are mapped by the AT-TLS policy rule to a policy with the SecondaryMap attribute specified as `On`. If entries are found using both methods, the one found by the AT-TLS policy rule is used unless the one found by the SecondaryMap value has a higher priority.

   If a TCP connection is not matching the expected rule, do one of the following:

   - Ensure that the AT-TLS policies are active and that no errors occurred. Message EZZ8438I is issued if Policy Agent encountered any errors while processing the AT-TLS policy. If errors occurred, review the Policy Agent logs for details on the error and correct the AT-TLS policy. You can use OBJERR to search the Policy Agent logs to find the errors.
   - Verify the rule and actions that the policy mapped to and the priority of the rule. You can use the **pasearch** command can be used to view the active AT-TLS policy. AT-TLS message EZD1281I is issued with all the parameters used to map to the AT-TLS policy, if trace level 4 is on.
4. If an error message was issued by AT-TLS, review the syslogd files for message EZD1286I or the TCP/IP joblog for message EZD1287I. The error message might provide information about correcting the problem.

5. If the error is recreatable, turn on an AT-TLS trace for the connection. Turn on the trace by coding a TTLSRule specific to the failing connection. Include a TTLSConnectionAction statement that has the Trace statement set to 255 (All). If configuring using the IBM Configuration Assistant for z/OS Communications Server, the trace level can be set in each Connectivity Rule.
6. If the problem cannot be resolved from the trace, perform a packet trace or a CTRACE with option TCP to provide additional debugging information and contact IBM service.

AT-TLS traces

By default, AT-TLS uses the syslog facility name daemon. Other TCP/IP functions, for example the SNMP agent, also specify the daemon facility name when writing records to syslogd. The job name and syslog facility name are the same. Filters cannot be used to direct the records to different output files. If you want AT-TLS records to go to a different output file, you can change the syslog facility name by configuring SyslogFacility Auth on the TTLSGroupAdvancedParms statement to direct the messages from that group to the Auth facility instead. You can then set up filtering based on the job name and facility in the syslogd configuration file to direct AT-TLS records to a different output file.

If you are configuring using the IBM Configuration Assistant for z/OS Communications Server, you can modify the syslog facility name from the AT-TLS: Image Level Settings panel.

AT-TLS traces are enabled by setting the AT-TLS policy statement Trace to a nonzero value. A Trace statement can be configured on a TTLSGroupAction, TTLSEnvironmentAction or TTLSConnectionAction statement. See the z/OS Communications Server: IP Configuration Reference for more details about AT-TLS policy statements. The Trace levels enable different AT-TLS messages to be issued. The sum of the numbers associated with each level of tracing that you want is the value that should be specified.

If you are configuring using the IBM Configuration Assistant for z/OS Communications Server, you can set the default trace level on the AT-TLS: Image Level Settings panel, and you can override the trace level for each Connectivity Rule.

Table 59 lists the trace level, the generated AT-TLS messages, and the syslog priority.

Table 59. AT-TLS trace levels

Trace level	Traced information	Syslog priority
1 Error (to Joblog)	EZD1287I	NA
2 Error	EZD1286I	err
4 - Info	EZD1281I, EZD1283I	info
8 - Event	EZD1282I, EZD1283I	debug
16 - Flow	EZD1282I, EZD1283I, EZD1284I	debug
32 - Data	EZD1285I	debug

Tip: Setting the Trace level to 6 enables both error messages and info messages.

The information messages trace when an AT-TLS connection is mapped to a policy (EZD1281I) and when the secure connection is successfully negotiated (EZD1283I), including the security protocol and cipher used. Using syslogd's filtering parameters, a separate log file could be kept for AT-TLS info and error messages, enabling AT-TLS connections to be tracked.

Tip: Trace level 32 shows all the SSL headers sent and received.

Each secure connection is uniquely identified by its connection ID (ConnID). You can use the ConnID to follow a connection through the AT-TLS trace.

Sample AT-TLS trace

Figure 92 on page 711 shows an example trace of a generic server processing a secure connection. The standard syslogd prefix information has been removed from the trace.

Trace level 255 was used to generate this trace.

```
11:10:25 TCPCS3   EZD1281I TTLS Map   CONNID: 00000025 LOCAL: 9.42.104.156..21 REMOTE: 9.27.154.171..1271
                           JOBNAME: FTPD2 USERID: FTPD TYPE: InBound STATUS: Enabled RULE: ftp_serv_21
                           ACTIONS: grp_act1 env_act_serv **N/A**   1
11:10:28 TCPCS3   EZD1283I TTLS Event GRPID: 00000001 ENVID: 00000000 CONNID: 00000025  RC:    0
  Connection Init
11:10:28 TCPCS3   EZD1282I TTLS Start GRPID: 00000001 ENVID: 00000001 CONNID: 00000000 Environment Create
                           ACTIONS: grp_act1 env_act_serv **N/A**   2
11:10:28 TCPCS3   EZD1283I TTLS Event GRPID: 00000001 ENVID: 00000002 CONNID: 00000000  RC:    0
  Environment Master
                           Create 00000001
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Call
                           GSK_ENVIRONMENT_OPEN - 7F1DB058
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_KEYRING_FILE -  FTPDsafkeyring    3
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_CLIENT_AUTH_TYPE -  FULL
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_SESSION_TYPE -  SERVER
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_PROTOCOL_SSLV2 -  ON
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_PROTOCOL_SSLV3 -  ON
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_PROTOCOL_TLSV1 -  ON
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_IO_CALLBACK -
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_SSL_HW_DETECT_MESSAGE -  1
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Call
                           GSK_ENVIRONMENT_INIT - 7F1DB058
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000002 CONNID: 00000025  RC:    0 Set
                           GSK_SSL_HW_DETECT_MESSAGE -  NULL
11:10:28 TCPCS3   EZD1283I TTLS Event GRPID: 00000001 ENVID: 00000002 CONNID: 00000000  RC:    0
  Environment Master
                           Init 7F1DB058
11:10:28 TCPCS3   EZD1283I TTLS Event GRPID: 00000001 ENVID: 00000001 CONNID: 00000000  RC:    0
  Environment
                           Link 7F1DB058 00000002
11:10:28 TCPCS3   EZD1282I TTLS Start GRPID: 00000001 ENVID: 00000001 CONNID: 00000025 Initial Handshake
                           ACTIONS: grp_act1 env_act_serv **N/A** HS-Server     4
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC:    0 Call
                           GSK_SECURE_SOCKET_OPEN - 7F0CA118
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC:    0 Set
                           GSK_FD - 00000025
11:10:28 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC:    0 Set
                           GSK_USER_DATA - 7F1DB330
```

```
11:10:28 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER 807A010301   5
11:10:28 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER
   0051000000200000040100800000050000 2F000033000032
   00000A0700C000001600001300000906004000001500001200000302008000000800001400001100000100000
   020000180000340000 1B00001A00001700001941E69D75F7DCB55234895D884B271253A522E4BE211250F546
                          4FE5C5AB980FBD
11:10:28 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 SEND CIPHER
   160301029002000046030141E69D753469372857A71168D9
   9D7B93AD6CC30F6F6BDF7F774929CD4D2E8E2A2000000026091B9AAB04F700000000000000000000000000000
   41E69D75000000010005000B00023E00023B000238308202343082019DA003020102020100300D06092A8648
   86F70D0101050500302E310B30090603550406130275733110300E060355040B130774657374696E67310D30
   0B0603550403130446545044301E170D303430383039303430305A170D3035303831303   6
11:10:28 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER 1603010086
11:10:28 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER
   10000082008037A6573A4C160A8C0810C542A1CEB73A9FF5
   899D767711EF3BF86D4C2D2743837AA4D5E247DE35F79C8A71A9E6A18DF8CC845D5E0F8F386DF84D746A4004
   B641C14DD7A002FAC5538ED52E3194C2ADE6010381BFC70D1CA6D9F34EDC0F345F0A015575A6C9D85602B1BF
   2877760BA91FC6296625A16A274426112C65DB7A2685
11:10:29 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER 1403010001
11:10:29 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER 01
11:10:29 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER 1603010024
11:10:29 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 RECV CIPHER
   789DBBACAE9D6F19F62B1AF2B529B1850F7057A6EDDE64CD 2301D91CA43C4EBBB5A3DFE5
11:10:29 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 SEND CIPHER 140301000101
11:10:29 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 SEND CIPHER
   603010024FE6548CCBA0D820D73FF439A6B475B4116BCE4
   6FF225DAE1A0F7EC2AEA4690595E63F036
11:10:29 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC:    0 Call
                          GSK_SECURE_SOCKET_INIT - 7F0CA118
11:10:29 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC:    0 Get
                          GSK_CONNECT_SEC_TYPE -  TLSV1
11:10:29 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC:    0 Get
                          GSK_CONNECT_CIPHER_SPEC -  05
11:10:29 TCPCS3   EZD1283I TTLS Event GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC:    0
   Initial Handshake
                          7F0CA118 7F1DB058 TLSV1 05  7
11:11:05 TCPCS3   EZD1285I TTLS Data  CONNID: 00000025 SEND CIPHER
   1503010016D47A7AEC70D317976ACEEF3418CDCC8B2DF7
   D3491D      8
11:11:13 TCPCS3   EZD1283I TTLS Event GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC: 0 Receive
   Reset
11:11:13 TCPCS3   EZD1282I TTLS Start GRPID: 00000001 ENVID: 00000001 CONNID: 00000025 Connection Close
                          ACTIONS: grp_act1 env_act_serv **N/A**    9
11:11:13 TCPCS3   EZD1284I TTLS Flow  GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC: 0 Call
                          GSK_SECURE_SOCKET_CLOSE - 7F0CA118
11:11:13 TCPCS3   EZD1283I TTLS Event GRPID: 00000001 ENVID: 00000001 CONNID: 00000025  RC: 0 Connection
                          Close 7F0CA118 7F1DB058
```

Figure 92. Example trace of a generic server processing

The following information corresponds to the line numbers in Figure 92.

1. A TCP connection has mapped to an AT-TLS rule. The parameters used to search the AT-TLS rules are listed. The TTLSRule, TTLSGroupAction, TTLSEnvironmentAction, and TTLSConnectionAction names are also displayed. Note the ConnID for the connection. This ConnID appears in all future AT-TLS messages for this connection.
2. AT-TLS is creating an environment instance for the application.
3. AT-TLS is establishing the parameters for this environment. These parameters are obtained from the TTLSEnvironmentAction statement. System SSL calls are made to set up the parameters. This trace message is defining the key ring to be used by this environment.

4. AT-TLS has successfully set up the secure environment and is now initializing the secure connection. This initiates network flows with the remote partner.
5. Secure data has been received for this connection. During secure handshake, all the data is traced. For this trace example, some of the data has been removed.
6. Secure data is being sent for this connection.
7. The secure handshake has completed. The protocol negotiated (TLSV1) and the cipher suite negotiated(05) are displayed.
8. AT-TLS is sending a secure alert message, because the application closed the socket.
9. The secure connection is being closed.

AT-TLS return codes

AT-TLS error message EZD1286I is issued to syslogd to report any errors that occur on an AT-TLS connection when the trace level 2 (Error) is set. AT-TLS error message EZD1287I is issued to the TCP/IP job log to report any errors that occur on an AT-TLS connection when the trace level 1 (Error) is set. These messages include the event that AT-TLS was processing and the return code that indicates a failure. Return codes 5001–5999 describe AT-TLS errors that can be corrected by the user. For more information about these return codes, see Table 61 on page 714. Return codes 6001–6999 describe internal AT-TLS errors. Contact IBM with the error message and syslog information, if available. Any other return code is defined by System SSL. For more information about these return codes, see z/OS Cryptographic Services System SSL Programming.

Table 60 lists some common System SSL return codes and possible causes.

Table 60. Common System SSL return codes

Return code	Event	Possible cause and solution
202	Environment Init	The key ring cannot be opened because the user does not have permission. Check the following items: • Look at message EZD1281 to verify the user ID being used for this connection and the TTLSEnvironmentAction statement that is mapped to this connection. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, you can specify the key ring on either the AT-TLS: Image Level Settings panel or on each Traffic Descriptor. • Ensure that the correct key ring is specified. • If you are using a RACF key ring, verify that all the steps in z/OS Communications Server: IP Configuration Guide are followed for this user ID.

Table 60. Common System SSL return codes (continued)

Return code	Event	Possible cause and solution
402	Connection Init	An SSL cipher suite could not be agreed upon between the client and server. Check the following items: • If V2Ciphers or V3Ciphers are coded, verify that the remote end supports at least one of the cipher suites coded. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, the ciphers are selected for each Security Level. • Verify that the certificate that is being used for the connection supports the cipher suites. For example, V3 Cipher suite TLS_DH_DSS_WITH_DES_CBC_SHA(0C) requires a certificate that is defined with a Diffie-Hellman key. • If ClientCurves is coded, verify that the server is using a certificate that supports one of the elliptic curves. • For ciphers defined as exportable, verify that the proper FMIDs to support the encryption level are installed.
406	Connection Init	An I/O error occurred on the socket. An I/O error occurs if the TCP socket is closed underneath the SSL protocol, such as when a reset is received. Check the following items: • Ensure that the remote partner is enabled for secure connections. • Determine whether the secure negotiation completed. Use the AT-TLS Data trace level to determine whether the secure negotiation completed. • Verify that the TCP data flows are sent by the remote partner. Use a TCP/IP packet trace to verify that the TCP data flows are sent by the remote partner.
412	Connection Init	A common SSL protocol type cannot be agreed upon by both partners. This disagreement occurs if both partners do not support the same SSL protocol, as when the client supports only SSLv2 and the server supports only TLSv1. AT-TLS supports only SSLv2, SSLv3, and TLSv1. Check the following items: • Determine the protocols that are supported by the remote partner. • Code a TTLSEnvironmentAdvancedParms statement, which enables the common protocols. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, use a Security Level with cipher levels supported by the remote partner.

Table 60. Common System SSL return codes (continued)

Return code	Event	Possible cause and solution
422	Connection Init	A v3Cipher that is not valid is found. Check the following items: • Determine whether the v3Cipher statement is coded. • Verify that the proper SSL FMIDs are installed to support the ciphers specified.
428	Connection Init	The private key cannot be obtained from the certificate. If the private key is stored in ICSF, ensure that ICSF was started before the Policy Agent starting. The TCP/IP job log contains a message that indicates whether System SSL detected ICSF services as available. If ICSF was not started before the Policy Agent, see technote 1433201 for instructions on how to update the AT-TLS policy to detect that ICSF is now available.
434	Connection Init	The certificate key is not compatible with the negotiated cipher suite. Ensure that the certificate that is being used supports the cipher suites that are coded with V2Ciphers or V3Ciphers. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, the ciphers are selected in each Security Level.

Table 61 lists some common AT-TLS return codes and possible causes.

Table 61. AT-TLS return codes

Return code	Possible cause and solution
5001	ClientAuthType is set to Required or SAFCheck, but the client did not provide a certificate. Verify that the client supports client authentication and is configured to send its certificate during secure negotiation.
5002	ClientAuthType is set to SAFCheck, but the certificate that is supplied by the client is not defined to SAF subsystem. If you are using RACF, define the client certificate with the RACDCERT command. For more information about using the RACDCERT command, see z/OS Security Server RACF Security Administrator's Guide.
5003	Clear text data is received on the connection from the remote partner instead of secure data. The connection is terminated. Check the following items: • Ensure that the remote client is enabled for secure connections. • If the policy is defined with ApplicationControlled `On`, ensure that the application read all the cleartext data before it started the secure handshake. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, the Application Controlled setting is done in each Traffic Descriptor.

Table 61. AT-TLS return codes (continued)

Return code	Possible cause and solution
5004	The first HandshakeTimeout interval expired without receiving secure data from the remote partner. The timer is set for the number of seconds specified by the HandshakeTimeout value when the secure connection is initiated. When the first secure data is received from the remote partner, the timer is canceled. Check the following items: • This return code can occur if both sides of the connection are configured to be the server in the secure handshake. Review the configuration to ensure that one side acts as the client. For AT-TLS, you can specify the HandshakeRole value in either the TTLSEnvironmentAction or the TTLSConnectionAction statement. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, configure the Handshake Role value in each Traffic Descriptor. • Increase the HandshakeTimeout value if the remote partner is not responding within the time interval. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, you can set the Timeout value in each Traffic Descriptor; you can override the value in each Connectivity Rule.
5005	The second HandshakeTimeout interval expired and the secure handshake is not finished. This interval is set to 10 times the HandshakeTimeout interval. The secure negotiation is started and the initial secure message is received from the remote partner. • If the remote partner is an interactive application, such as requiring the user to select a certificate, either increase the HandshakeTimeout value or have the user try the connection again. • The HandshakeTimeout value might need to be increased if LDAP is being used to manage certificates. Increasing the value provides more time for the LDAP processing to occur. If you are configuring by using the z/OS Network Configuration Assistant, the Handshake Timeout value can be set in each Traffic Descriptor and can be overridden in each Connectivity Rule.

Table 61. AT-TLS return codes (continued)

Return code	Possible cause and solution
5006	The connection is using a TTLSEnvironmentAction statement that failed to initialize a System SSL environment. • Use the syslog to determine why the System SSL environment failed to initialize. • If the TTLSEnvironmentAction statement is in error, make the necessary corrections. A System SSL environment is initialized for the corrected TTLSEnvironmentAction statement and new connections use that environment. • If a SAF configuration change is needed (such as changing a certificate in the key ring), make the change and then update the EnvironmentUserInstance parameter in the TTLSEnvironmentAction statement to reflect a changed action. A System SSL environment is initialized by using the modified RACF configuration and new connections use that environment. If you are configuring by using the z/OS Network Configuration Assistant to pick up changes that are made to a key ring, go to the AT-TLS Image Level Settings panel and click the **Reaccess Key Rings** button and update the Instance ID for the changed key ring.
5007	Application data is read during processing of ciphertext negotiation. Collect the syslogd output or job log output and contact IBM.
5008	Application data was received after the local application closed the TCP connection. The data could not be presented to the application. • Review the local and remote applications to ensure that the TCP sockets are being closed correctly in the application flow. • If further diagnostic information is needed, set the trace level to 255, to trace the data flow and AT-TLS processing.
5009	AT-TLS was unable to obtain TCPIP private storage. Obtain a console dump of TCPIP and contact IBM
5010	AT-TLS was unable to obtain the ACEE for an application. Save the syslogd output and contact IBM
5011	AT-TLS does not have an Envar object for the applications ACEE. Save the syslogd output and contact IBM
5012	An internal AT-TLS error occurred. Save the syslogd output and contact IBM
5013	AT-TLS was unable to clone the SAF environment for the application. Save the syslogd output and contact IBM.
5014	AT-TLS was unable to extract ACEE into ENVAR value. Save the syslogd output and contact IBM.
5015	AT-TLS was unable to process the connection because the connection is already terminated. Review the syslogd output to determine whether the connection is terminated by the remote partner. TTLS trace level 8 (flow) and 16 (event) can be used to gather more information.
5016	AT-TLS attempted to read ciphertext negotiation data, but an internal error occurred. Save the syslogd output and contact IBM

Table 61. AT-TLS return codes (continued)

Return code	Possible cause and solution
5017	The application tried to write data on a secure connection that is closed by the remote application. • Review the local and remote applications to ensure that the TCP sockets are being closed correctly in the application flow. • If further diagnostic information is needed, set the trace level to 255, to trace the data flow and AT-TLS processing.
5018	An internal error occurred processing a TTLSGroupAction. Save the syslogd output and contact IBM.
5019	Task level security could not be created. BPX1TLS failed. Save the syslogd output and contact IBM.
5020	AT-TLS was unable to load the GSKSSL library. Ensure that the SIEALNKE PDSE library is available to the TCPIP started task. For more information, see z/OS Cryptographic Services System SSL Programming.
5021	The HandshakeTimeout interval expired for the SIOCTTLSCTL TTLS_Stop_Connection request without receiving a close notify alert from the remote peer. The timer is set for the number of seconds that the HandshakeTimeout value specifies when the TTLS_Stop_Connection request is initiated. The timer is canceled when a close notify alert is received from the remote peer. Increase the HandshakeTimeout value if the remote peer is not responding within the time interval. If you are configuring by using the IBM Configuration Assistant for z/OS Communications Server, you can set the timeout value in each Traffic Descriptor; you can override the value in each Connectivity Rule.
5022	Encrypted application data is received from the remote peer when the SIOCTTLSCTL TTLS_Stop_Connection request is being processed. All application data that needs to be encrypted must be sent before the TTLS_Stop_Connection request. The application protocol needs to ensure all sending and receiving of secure data on the connection is complete before TTLS_Stop_Connection is requested. Review the application protocol to determine why the secure data is sent on the connection.
5023	AT-TLS called initACEE with a nested ENVR object and requested a managed ACEE, which is not supported. If AT-TLS was processing a data connection from the FTP server, ensure the AT-TLS policy has SecondaryMap On coded for the FTP control connection. A separate TTLSRule for the FTP data connection is not supported. Otherwise, save the syslogd output and contact IBM.
5024	AT-TLS was unable to enable FIPS 140 support. See message EZD2026I for more details about the error that is received from System SSL.
Return codes 6001–6999 describe internal AT-TLS errors.	An internal AT-TLS error occurred. Contact IBM with the error message and syslog information, if available.

SIOCTTLSCTL ioctl return codes

The SIOCTTLSCTL ioctl provides the interface for an application to query and control AT-TLS. Table 62 describes the error codes that can be returned on this ioctl, along with the conditions under which each can occur. Also included for each is an indication of whether the query data fields in the ioctl contain valid returned data.

Table 62. SIOCTTLSCTL error codes

Errno (1)	Errnojr	IOCTL request specified (2)	Condition causing Error	Valid Data? (3)
EAcces	JrConnDeniedPolicy	`INIT_CONNECTION, RESET_SESSION, RESET_CIPHER, STOP_CONNECTION`	Mapped policy indicates that the application cannot request AT-TLS security for the connection (ApplicationControlled Off)	Yes
EAlready	JrAlreadyActive	`INIT_CONNECTION, STOP_CONNECTION`	An INIT_CONNECTION or STOP_CONNECTION request is previously received for the connection	Yes
EConnReset	JrTTLSHandshakeFailed	Any	Initial handshake was in progress and socket is a blocking socket. Request blocked for handshake to complete. Handshake failed.	No
EInProgress	JrOK	`INIT_CONNECTION, STOP_CONNECTION`	Initial handshake or stop secure connection is started and socket is a non-blocking socket. (4)	Yes
EInval	JrInvalidVersion	Any	Bad ioctl version number specified.	No
EInval	JrSocketCallParmError	Any	Length of input data is not length of ioctl structure.	No
EInval	JrSocketCallParmError	Not valid	Request type specified is not valid.	No
EInval	JrSocketCallParmError	RETURN_CERTIFICATE	Certificate buffer pointer = 0 or certificate buffer length = 0.	No
EInval	JrSocketCallParmError	! RETURN_CERTIFICATE	Certificate buffer pointer != 0 or certificate buffer length != 0 and TTLS_Version is 1.	No
EMVSErr	JrUnexpectedErr	Any	Policy was not mapped before ioctl call and an error is encountered upon policy map during ioctl call.	No
ENoBufs	JrBuffTooSmall	RETURN_CERTIFICATE	The certificate buffer provided is too small.	Yes (5)
ENoBufs	JrBuffTooSmall	QUERY_ONLY	A TTLS_Version 2 request was issued, but the buffer was too small.	Yes (6)

Table 62. SIOCTTLSCTL error codes (continued)

Errno (1)	Errnojr	IOCTL request specified (2)	Condition causing Error	Valid Data? (3)
ENotConn	JrGetConnError	Any	TCP connection is not yet in established state or is reset.	No
EOpNotSupp	JrOptNotSupported	INIT_CONNECTION, RESET_SESSION, RESET_CIPHER STOP_CONNECTION	Mapped policy indicates that AT-TLS is not enabled for the connection (TTLSEnabled Off).	Yes
EPerm	JrSocketCallParmError	INIT_CONNECTION with RESET_SESSION or RESET_CIPHER or STOP_CONNECTION, STOP_CONNECTION with RESET_SESSION or RESET_CIPHER, ALLOW_HSTIMEOUT without INIT_CONNECTION	Combination of requests specified is not permitted.	No
EPipe	JrUnexpectedErr	`INIT_CONNECTION`, `RESET_CIPHER`, `STOP_CONNECTION`	TCP connection is no longer in Established state. Two-way communication is not possible.	Yes
EProto	JrGetConnErr	RESET_SESSION, RESET_CIPHER	An INIT_CONNECTION request is not received for the connection.	Yes
EProto	JrInvalidVersion	`RESET_CIPHER`, `STOP_CONNECTION`	Connection is secured by using SSL version 2.	Yes
EProto	JrConnDeniedPolicy	ALLOW_HSTIMEOUT	The `TTLS_ALLOW_HSTIMEOUT` option was requested but the HandshakeRole is a client or the HandshakeTimeout value is 0.	Yes
EProto	JrTTLSStopReadDataPending	STOP_CONNECTION	A STOP_CONNECTION request is issued on a connection that has unread application data.	Yes
EProto	JrTTLSStopWriteDataPending	STOP_CONNECTION	A STOP_CONNECTION request is issued on a connection that has outstanding application write requests.	Yes
EProtoType	JrSocketTypeNotSupported	Any	Socket is not a TCP socket.	No
EWouldBlock	JrOK	Any	SSL handshake is in progress and socket is a non-blocking socket. (4)	Yes

Table 62. SIOCTTLSCTL error codes (continued)

Errno (1)	Errnojr	IOCTL request specified (2)	Condition causing Error	Valid Data? (3)

Notes:

1. For the numeric values of the error codes, see the following topics:
   - If you are using the BPX APIs, see the return codes in z/OS UNIX System Services Messages and Codes.
   - If you are using the TCP/IP APIs, see the return codes in z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.
2. The entry `Any` indicates that any valid request or valid combination of request types was specified as follows:

   request_type
   The listed request_type value was specified alone or in any valid combination of request_type.

   request_type, request_type[, request_type]
   One of the listed request types was specified alone or in any valid combination of request types.

   request_type with request_type
   The listed pair of request types was specified together.

   ! request_type
   Any valid combination of request types that does not include the listed request_type was specified.
3. `Yes` indicates that query data fields in the ioctl control block contain valid returned data. `No` indicates that the query data fields are unmodified.
4. For a non-blocking socket, you can wait for the handshake to complete by issuing Select or Poll for Socket Writable.
5. Certificate is not returned because the buffer was not large enough to hold it.
6. Output data is returned for output requests which completely fit in the buffer provided.

Chapter 29. Diagnosing Defense Manager daemon problems

This topic describes how to diagnose Defense Manager daemon (DMD) problems, and contains the following subtopics:

- "Overview of diagnosing Defense Manager daemon problems"
- "Defense Manager daemon debug information" on page 727
- "TCP/IP services component trace for the Defense Manager daemon" on page 728
- "Enabling CTRACE at Defense Manager daemon startup" on page 729

Overview of diagnosing Defense Manager daemon problems

The DMD oversees the addition, modification and deletion of TCP/IP stack defensive filters. Problems with the DMD may be categorized as follows:

- DMD configuration problems
- DMD internal problems
- DMD problems interacting with an external component such as:
  - **ipsec** command
  - Secure Access Facility

The DMD provides log output using syslogd and internal trace information using component trace (CTRACE). The log output is sufficient for diagnosing most DMD problems and is the first place to look if you suspect a problem.

Table 63 on page 722 lists common DMD problems.

© Copyright IBM Corp. 2000, 2015

Table 63. Common defense manager daemon (DMD) problems

Problem	Symptom	Cause/response

ipsec command cannot be used to update or delete a defensive filter installed in the stack	The **ipsec -F display** command lists defensive filters that cannot be updated or deleted. An **ipsec -F update** command to update a defensive filter fails with message:
	`EZD1550I Defense Manager daemon reported an error - filter is not found`
	
	An **ipsec -F delete** command to delete a defensive filter fails with message:
	`EZD1546I Defensive filter filtername was not found in stack stackname`
		Under normal circumstances the DMD's defensive filter files are in sync with the defensive filters in each running stack. However, there are several operational errors that can lead to a mismatch where defensive filters are installed in the stack but the DMD does not have any knowledge of them. The **ipsec –F display** command will display the defensive filters installed in the stack. However, the filters cannot be referenced by name to update or delete them.
		
		One possible cause of this mismatch is if the administrator removes the DmStackConfig statement from the DMD configuration while there are active defensive filters in the stack. If this was done to disable defensive filtering then only a partial disablement was done. No new defensive filters will be added to the stack but the existing ones remain. The user has 2 choices to complete the disablement. The MODIFY FORCE_INACTIVE command can be used to disable defensive filtering. Or the DmStackConfig statement can be added back to the configuration file with Mode Inactive. In both cases, any existing defensive filters will be deleted from the stack.
		
		If the DmStackConfig statement was removed inadvertently and defensive filtering should remain enabled for the stack, add the DmStackConfig statement back to the configuration file and refresh the configuration. All of the active defensive filters should be addressable again and new defensive filters can be added.
		
		Refreshing the DMD configuration with the wrong configuration file could result in the apparent deletion of a stack. For example, if DMD is started with configuration file /etc/security/dmd.conf which has a DmStackConfig statement for TCPCS2 and TCPCS3, then a refresh is done with configuration file /etc/security/dmd.conf1 which only has a DmStackConfig statement for TCPCS2, it will appear that DmStackConfig has been deleted for TCPCS3. The configuration should be refreshed with the correct configuration file. Any active defensive filters should be addressable again.
		
		Another possible cause is that the DMD is started with the DefensiveFilterDirectory parameter in the DMD configuration file set to the wrong location. To correct the problem, stop the DMD. The DefensiveFilterDirectory parameter cannot be changed on a MODIFY REFRESH. Correct the DefensiveFilterDirectory value and restart the DMD. All of the active defensive filters should be addressable again.

Table 63. Common defense manager daemon (DMD) problems (continued)

Problem	Symptom	Cause/response

The DMD cannot be started. The configuration file cannot be read.	The following message sequence is written to the MVS console:
	`EZD1610I THE DEFENSE MANAGER DAEMON INITIALIZATION SEQUENCE HAS BEGUN`
	`EZD1617I AN ERROR OCCURRED WHILE READING THE DEFENSE MANAGER DAEMON CONFIGURATION FILE /etc/security/dmd.conf - RETURN CODE 2`
	`EZD1604I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS BEGUN`
	`EZD1605I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS COMPLETED`
		The configuration file must exist and the DMD must have the right permissions to read it.
		If the configuration data is not in the default location (that is, /etc/security/dmd.conf), ensure that the environment variable DMD_FILE is set to the name of the MVS data set or z/OS UNIX file that contains the configuration data.
		If the DMD is defined with a non-zero UID, ensure that the DMD has permission to read the configuration file. This requires that the DMD user has both read access to the configuration file, as well as access to the directory containing the configuration file.
		Tip: The /var/dm directory must be set up to allow DMD to create, delete, read, and write files to it. You might want to create the configuration file in the /var/dm directory and use the DMD_FILE environment variable to specify the configuration file.

The DMD cannot be started. The /var/dm directory does not exist and cannot be created.	The following message sequence is written to the MVS console:
	`EZD1610I THE DEFENSE MANAGER DAEMON INITIALIZATION SEQUENCE HAS BEGUN`
	`EZD1604I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS BEGUN`
	`EZD1605I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS COMPLETED`
	The following message is written to syslog:
	`EZD1624I The Defense Manager daemon socket directory /var/dm does not exist and cannot be created - errno 111 EDC5111I Permission denied`
		The directory /var/dm must exist or the DMD must have the right permissions to create it. If the DMD needs to create the dm subdirectory, /var must exist.
		If the DMD is defined with a non-zero UID, create the /var/dm directory before starting the DMD. The directory should be owned by the DMD userid and the DMD should be able to create, delete, read, and write files to the directory.

Table 63. Common defense manager daemon (DMD) problems (continued)

Problem	Symptom	Cause/response

The DMD cannot be started, The directory /var/dm exists but the DMD is not able to create a file in the directory.	The following message sequence is written to the MVS console:
	`EZD1610I THE DEFENSE MANAGER DAEMON INITIALIZATION SEQUENCE HAS BEGUN`
	
	`EZD1732I AN ERROR OCCURRED WHILE TRYING TO ACCESS THE DMD DIRECTORY /var/dm - ERRNO 111 EDC5111I Permission denied.`
	
	`EZD1617I AN ERROR OCCURRED WHILE READING THE DEFENSE MANAGER DAEMON CONFIGURATION FILE /etc/security/dmd.conf - RETURN CODE 3`
	
	`EZD1604I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS BEGUN`
	`EZD1605I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS COMPLETED`
		If the DMD is defined with a non-zero UID, the directory should be owned by the DMD userid and the DMD should be able to create, delete, read, and write files to the directory.

Table 63. Common defense manager daemon (DMD) problems (continued)

Problem	Symptom	Cause/response

The DMD cannot be started. The defensive filter directory does not exist.	The following message sequence is written to the MVS console:
	`EZD1610I THE DEFENSE MANAGER DAEMON INITIALIZATION SEQUENCE HAS BEGUN`
	
	`EZD1621I AN ERROR OCCURRED WHILE TRYING TO ACCESS DEFENSIVE FILTER DIRECTORY /var/dm/filters - ERRNO 129 EDC5129I No such file or directory.`
	
	`EZD1617I AN ERROR OCCURRED WHILE READING THE DEFENSE MANAGER DAEMON CONFIGURATION FILE /etc/security/dmd.conf - RETURN CODE 3`
	
	`EZD1604I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS BEGUN`
	
	`EZD1605I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS COMPLETED`
		The defensive filter directory specified in the DMD configuration file on the DefensiveFilterDirectory keyword must be created before starting the DMD. The DMD must be able to create, delete, read, and write files to the directory.
		
		If the DMD is defined with a non-zero UID, the directory should be owned by the DMD userid and the DMD should be able to create, delete, read, and write files to the directory.

Table 63. Common defense manager daemon (DMD) problems (continued)

Problem	Symptom	Cause/response

The DMD cannot be started. Files cannot be written to the defensive filter directory.	The following message sequence is written to the MVS console:
	`EZD1610I THE DEFENSE MANAGER DAEMON INITIALIZATION SEQUENCE HAS BEGUN`
	
	`EZD1621I AN ERROR OCCURRED WHILE TRYING TO ACCESS DEFENSIVE FILTER DIRECTORY /var/dm/filters - ERRNO 111 EDC5111I Permission denied.`
	
	`EZD1617I AN ERROR OCCURRED WHILE READING THE DEFENSE MANAGER DAEMON CONFIGURATION FILE /etc/security/dmd.conf - RETURN CODE 3`
	
	`EZD1604I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS BEGUN`
	
	`EZD1605I THE DEFENSE MANAGER DAEMON SHUTDOWN SEQUENCE HAS COMPLETED`
		The DMD must be able to create, read, write, and delete files from the defensive filter directory specified in the DMD configuration file on the DefensiveFilterDirectory keyword.
		
		If the DMD is defined with a non-zero UID, the directory should be owned by the DMD userid and the DMD should be able to create, delete, read, and write files to the directory.

Table 63. Common defense manager daemon (DMD) problems (continued)

Problem	Symptom	Cause/response

The DMD starts but the process ID is not written to a file.	The following message is written to the MVS console if the directory portion of the PID file location does not exist:
	`EZD1603I THE DEFENSE MANAGER DAEMON FAILED TO WRITE ITS PROCESS ID 67108892 TO /var/dm/dmd.pid - ERRNO 129 ERRNO DESCRIPTION EDC5129I No such file or directory.`
	
	The following message is written to the MVS console if the DMD does not have permission to write the PID file to the directory:
	`EZD1603I THE DEFENSE MANAGER DAEMON FAILED TO WRITE ITS PROCESS ID 29 TO /var/dm/dmd.pid - ERRNO 111 ERRNO DESCRIPTION EDC5111I Permission denied.`
		The directory portion of the PID file location specified by the DMD_PIDFILE environment variable or defaulted to /var/dm must be created before starting the DMD. The DMD must be able to write the PID file to the directory.
		
		If the directory does not exist, create it. If the DMD is defined with a non-zero UID, the directory should be owned by the DMD userid and the DMD should be able to create, delete, read, and write files to the directory.

Defense Manager daemon debug information

The SyslogLevel parameter in the DMD configuration file controls the level of DMD internal debug information that is sent to syslog. See z/OS Communications Server: IP Configuration Reference for more information.

Abends during DMD processing

Messages and error-related information should be sent to the system console when an abend occurs during DMD processing. A dump of the error is needed unless the symptoms match a known problem. System dumps of the DMD include Language Environment data. The Language Environment IPCS verbexit LEDATA can be used to format this information. See z/OS Language Environment Debugging Guide for more information. The following command is a sample IPCS **verbexit ledata** command:

```
verbx ledata 'asid(68) tcb(007E5E88) ceedump nthreads(*)'
```

Tip: In this example, the DMD asid is 0x68 and the address of the abended DMD TCB is 0x007E5E88.

DMD error codes

Several messages display a return code and reason generated by the DMD. These return codes and reasons are displayed by the **ipsec** command.

TCP/IP services component trace for the Defense Manager daemon

The DMD uses component trace support to trace internal operations. For short descriptions of other tracing procedures, such as displaying trace status, see Chapter 5, "TCP/IP services traces and IPCS support," on page 45.

For detailed information, see the following:

- z/OS MVS Diagnosis: Tools and Service Aids for information about component trace procedures.
- z/OS MVS Initialization and Tuning Reference for information about the component trace SYS1.PARMLIB member.
- z/OS MVS System Commands for information about commands.
- z/OS MVS Programming: Authorized Assembler Services Guide for procedures and return codes for component trace macros.

CTRACE options

You can specify component trace options at DMD initialization or after the DMD has initialized.

Table 64 lists the DMD trace options.

Table 64. Defense manager daemon (DMD) trace options

Trace event	Description
ALL	Select all types of records. Be aware that this option slows performance.
MINIMUM	Select the DMD's minimum level of tracing. This level includes the INIT, EXCEPT, and TERM categories.
INIT	Select the DMD initialization information.
TERM	Select the DMD termination information.
EXCEPT	Select the DMD exception information.
CONFIG	Select the DMD configuration information.
COMMANDS	Select processing of DMD commands from the console or command line.
LOGMSGS	Select the DMD syslog messages. These entries can be used to easily correlate system log messages to a specific point in the CTRACE log.
ROUTING	Select the DMD threading and request dispatching information.
SERIAL	Select the DMD serialization information.
EVENT	Select the DMD event information.
SOCKETS	Select the DMD socket information.
PERFORM	Select the DMD performance information.
REQUESTS	Select the DMD request/response information.
FLOW	Select the DMD code flow information.
STORAGE	Select the DMD storage information.
CONTROL	Select the DMD control information.
DEBUG	Select the DMD debugging information.
VERBOSE	Select the DMD verbose debugging information.

Enabling CTRACE at Defense Manager daemon startup

A default minimum component trace is always started during DMD initialization. Use a parmlib member to customize the parameters that are used to initialize the trace. The default DMD component trace parmlib member is the SYS1.PARMLIB member CTIDMD00. The parmlib member name can be changed using the DMD_CTRACE_MEMBER environment variable.

Rule: The DMD reads the DMD_CTRACE_MEMBER environment variable only during initialization. Changes to DMD_CTRACE_MEMBER after server initialization have no affect.

Restriction: In addition to specifying the trace options, you can also change the DMD trace buffer size. The buffer size can be changed only at DMD initialization and has a maximum size of 256 MB.

If the CTIDMD00 member or the member that is specified in DMD_CTRACE_MEMBER is not found when starting the DMD, the following message is issued:

```
IEE5381 memberName MEMBER NOT FOUND IN PARMLIB
```

When this occurs, the DMD component trace is started with a buffer size of 1 MB and the MINIMUM tracing option.

```
/*********************************************************************/
/* z/OS Communications Server                                        */
/* SMP/E Distribution Name: CTIDMD00                                 */
/*                                                                   */
/* PART Name: CTIDMD00                                               */
/*                                                                   */
/* Copyright:                                                        */
/* Licensed Materials - Property of IBM                              */
/* 5694-A01                                                          */
/* Copyright IBM Corp. 2008                                          */
/* Status:  CSV1R10                                                  */
/*                                                                   */
/* DESCRIPTION = This parmlib member causes component trace for      */
/* the TCP/IP Defense Manager daemon application to be initialized   */
/* with a trace buffer size of 1M                                    */
/*                                                                   */
/* This parmlib member only lists those TRACEOPTS                    */
/* values specific to DM. For a complete list                        */
/* of TRACEOPTS keywords and their values see                        */
/* z/OS MVS Initialization and Tuning Reference.                     */
/*                                                                   */
/*********************************************************************/
 TRACEOPTS
 /* ---------------------------------------------------------------- */
 /* Optionally start external writer in this file (use both          */
 /* WTRSTART and WTR with same wtr_procedure)                        */
 /* ---------------------------------------------------------------- */
 /* WTRSTART(wtr_procedure)                                          */
 /* ---------------------------------------------------------------- */
 /* ON OR OFF: PICK 1                                                */
 /* ---------------------------------------------------------------- */
 ON
 /* OFF                                                              */
 /* ---------------------------------------------------------------- */
 /* BUFSIZE: A VALUE IN THE RANGE OF 128K TO 256M                    */
 /* CTRACE buffers reside in the Defense Manger daemon's private     */
 /* storage which is in the region's address space.                  */
 /* ---------------------------------------------------------------- */
 BUFSIZE(1M)
 /* WTR(wtr_procedure)                                               */
 /* ---------------------------------------------------------------- */
 /* OPTIONS: NAMES OF FUNCTIONS TO BE TRACED, OR "ALL"               */
 /* ---------------------------------------------------------------- */
 /* OPTIONS(          */
 /*   'ALL     '      */
 /* , 'MINIMUM '      */
 /* , 'INIT    '      */
 /* , 'TERM    '      */
 /* , 'EXCEPT  '      */
 /* , 'CONFIG  '      */
 /* , 'ROUTING '      */
 /* , 'COMMANDS'      */
 /* , 'LOGMSGS '      */
 /* , 'SERIAL  '      */
 /* , 'EVENT   '      */
 /* , 'SOCKETS '      */
 /* , 'REQUESTS'      */
 /* , 'FLOW    '      */
 /* , 'STORAGE '      */
 /* , 'CONTROL '      */
 /* , 'VERBOSE '      */
 /* , 'DEBUG   '      */
 /* , 'PERFORM '      */
 /* )                 */
```

Figure 93. SYS1.PARMLIB member CTIDMD00

Steps for enabling the CTRACE at Defense Manager daemon startup

Edit the CTIDMD00 parmlib member appropriately and start the DMD.

Procedure

1. Edit the CTIDMD00 parmlib member and specify TRACEOPTS ON, the desired buffer size with the BUFSIZE() parameter and the desired CTRACE options. To direct the CTRACE to an external writer, also specify the name of the writer JCL procedure in the WTR() parameter. See Figure 93 on page 730.
2. Start the DMD.

Steps for disabling the CTRACE at Defense Manager daemon startup

Edit the CTIDMD00 parmlib member appropriately and start the DMD.

Procedure

1. To disable the CTRACE at DMD startup, edit the CTIDMD00 parmlib member and specify TRACEOPTS OFF.
2. Start the DMD.

Step for enabling the CTRACE after the Defense Manager daemon has started

After the DMD is started, enable the CTRACE to an internal buffer or to an external writer.

Procedure

Perform the appropriate step:

- Issue the following console commands to enable the CTRACE to an internal buffer:

```
TRACE CT,ON,COMP=SYSTCPDM,SUB=(dmd_jobname)
R xx,OPTIONS=(option[,option2...]),END
```

- Issue the following console commands to enable the CTRACE to an external writer:

```
TRACE CT,WTRSTART=writer_proc
TRACE CT,ON,COMP=SYSTCPDM,SUB=(dm_jobname)
R xx,OPTIONS=(option[,option2...]),WTR=writer_proc,END
```

Step for disabling the CTRACE after the Defense Manager daemon is started

After the DMD is started, disable the CTRACE to an internal buffer or to an external writer.

Procedure

Perform the appropriate step:

- Issue the following console commands to disable the CTRACE to an internal buffer:

```
TRACE CT,OFF,COMP=SYSTCPDM,SUB=(dm_jobname)
```

- Issue the following console commands to disable a CTRACE to an external writer:

```
TRACE CT,OFF,COMP=SYSTCPDM,SUB=(dm_jobname)
TRACE CT,WTRSTOP=writer_proc
```

Displaying the CTRACE status

To display the CTRACE status, issue the following console command:

```
D TRACE,COMP=SYSTCPDM,SUB=(dm_jobname)
```

Enabling CTRACE after Defense Manager daemon initialization

After DMD initialization, you must use the TRACE CT command to change the component trace options. Each time a new Component Trace is initiated, all prior trace options are turned OFF and the new options are put into effect. You can specify the trace options with or without the parmlib member. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for more information.

Formatting Defense Manager daemon trace records

You can format component trace records using IPCS panels or a combination of the IPCS panels and the CTRACE command, either from a dump or from external-writer files. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for details.

Enter any combination of values as options to filter the CTRACE entries. The options must be entered using the following format: `TYPE(option[,option]...)`.

You can use any of the options listed in Table 64 on page 728, except ALL and MINIMUM.

Chapter 30. Diagnosing IP security and defensive filter problems

This topic describes how to diagnose IP security problems including problems with defensive filters. It contains the following subtopics:

- "Overview of diagnosing IP security and defensive filter problems"
- "Steps for diagnosing IP security problems" on page 734
- "Steps for diagnosing defensive filter problems" on page 735
- "Steps for diagnosing the cause for missing IP security or defensive filter syslogd output" on page 737
- "Steps for verifying IP security and defensive filter operation" on page 741
- "Tools for diagnosing IP security and defensive filter problems" on page 753

Overview of diagnosing IP security and defensive filter problems

IPSec configuration files are input to the Policy Agent to establish a TCP/IP stack IP filter policy, Key Exchange policy, and LocalDynVpn policy. These configuration files consist of a number of configuration statements and parameters documented in z/OS Communications Server: IP Configuration Reference and can be configured manually into a flat file. Optionally, IBM provides an IBM Configuration Assistant for z/OS Communications Server, which provides wizards and a set of reusable objects (at a different level of abstraction than if configured manually). The IBM Configuration Assistant for z/OS Communications Server ultimately produces the Policy Agent configuration files on your behalf.

When diagnosing problems, it might be helpful to understand the relationship of the GUI level objects to the configuration file objects. Table 65 provides a brief mapping of these objects.

Table 65. GUI-level object mapping

Policy Agent Object	**IBM Configuration Assistant for z/OS Communications Server Object**
IpDataOffer	Configured in security levels implementing dynamic tunnels
IpDynVpnAction	Security level implementing dynamic tunnels A numeric suffix is appended to the Security Level name to guarantee uniqueness.
IpFilterRule	Connectivity rule A numeric suffix is appended to the connectivity rule name to guarantee uniqueness.
IpManVpnAction	Security level implementing manual tunnels A numeric suffix is appended to the security level name to guarantee uniqueness.
IpService	Configured in traffic descriptors A numeric suffix is appended to the traffic descriptor name to guarantee uniqueness.

© Copyright IBM Corp. 2000, 2015

Table 65. GUI-level object mapping (continued)

Policy Agent Object	IBM Configuration Assistant for z/OS Communications Server Object
IpTimeCondition	Defined within either Connectivity Rules or Security Levels implementing Manual Tunnels
KeyExchangeAction	Configured in connectivity rules A numeric suffix is appended to the connectivity rule name to guarantee uniqueness.
KeyExchangeRule	Configured in Connectivity Rules A numeric suffix is appended to the Connectivity Rule name to guarantee uniqueness.
LocalDynVpnRule	Configured in connectivity rules Names are user specified.

The Policy Agent installs IP security policy into the stack and the IKE daemon. Specifically, IP filter policy is installed in the stack and Key Exchange policy and LocalDynVpn policy are installed in the IKE daemon. The stack enforces IP filter policy after it has been successfully installed. The IKE daemon enforces Key Exchange policy and LocalDynVpn policy after they have been successfully installed. The Traffic Regulation Management daemon (TRMD) reports IP security events to syslogd on behalf of the stack.

Defensive filters:

Defensive filters are deny filters that can be added through the **ipsec** command, typically by an external security information and event manager that detects an attack. Defensive filters can only be installed in a TCP/IP stack that has IP security enabled. Defensive filters are given higher priority than IP security filters. That is, IP filter processing will first check a packet against any installed defensive filters for a match, before checking the IP security filters.

Problems can occur in the following areas:

- IP security policy installation
- IP security and defensive filter output to syslogd
- IP security operation
- Adding and managing defensive filters

Steps for diagnosing IP security problems

Diagnose IP security problems.

Procedure

Perform the following steps:

1. Issue **pasearch -v a** to see all IP security policies that are active in policy agent. See z/OS Communications Server: IP System Administrator's Commands for more information about the **pasearch -v a** command. If you are running multiple stacks, ensure that pasearch is reporting on the stack you are interested in. See Chapter 25, “Diagnosing Policy Agent problems,” on page 657 if you do not see the IP security policies that you expected

Tips:

- IP security policies that are active in the Policy Agent might not be active in the stack. Issue **ipsec -f display** and locate the Source field to determine the source of the policy that is active in the stack. If the Source field indicates `Stack Policy`, then the policy that is active in the Policy Agent corresponds to the policy that is active in the stack.
- Defensive filters are not defined in the policy agent configuration file so defensive filters are not displayed by **pasearch**.

2. Issue **ipsec -f display** to see how the stack mapped your IpFilterPolicy statement. See z/OS Communications Server: IP System Administrator's Commands for more information about the **ipsec -f** command. If you are running multiple stacks, ensure that your resolver configuration correctly identifies the stack you are interested in. Ensure that your IP security policies are correctly defined. See the IP security information in z/OS Communications Server: IP Configuration Guide.

   Tip: When the command **ipsec -f display** command is issued with a scope of **-c current**, any defensive filters installed in the stack will be displayed along with IP security filters.

Steps for diagnosing defensive filter problems

Diagnose defensive filter problems.

Procedure

Perform the following step to display defensive filters:

1. Issue **ipsec -F display** to display active defensive filters. See z/OS Communications Server: IP System Administrator's Commands for more information about the **ipsec -F** command. The defensive filters installed in a stack can be displayed or the global defensive filters can be displayed. If you are running multiple stacks and you want to display defensive filters installed in a specific stack, specify the **-p** *stacknameoption* or ensure that your resolver configuration correctly identifies the stack you are interested in. To display global defensive filters, specify the **-G** option.

Perform the following steps to determine why defensive filters are not being successfully added to a stack:

2. Ensure that IP security is enabled for the stack. Specify IPSECURITY on the IPCONFIG statement in the TCP/IP profile. In addition, specify IPSECURITY on the IPCONFIG6 statement in the TCP/IP profile if support is needed for IPv6 defensive filters. See z/OS Communications Server: IP Configuration Reference for more information about the IPCONFIG IPSECURITY and IPCONFIG6 IPSECURITY statements.
3. Ensure that the Defense Manager daemon (DMD) is managing defensive filters for the stack. The TCP/IP stack name must be listed in the DMD configuration file to enable defensive filters for the stack. The mode specified on the DmStackConfig statement for the stack must be Active or Simulate. See z/OS Communications Server: IP Configuration Reference for more information about the DMD configuration file.
4. Ensure that the user has security product authorization to issue the **ipsec** command to add a defensive filter. See "**ipsec** command security" in z/OS Communications Server: IP System Administrator's Commands for more

information about defining the necessary SERVAUTH profiles.

Perform the following step if administrative access is being denied by a defensive filter:

5. Exclude the administrator's IP address from defensive filter checking. Use the Exclude keyword on the DmStackConfig statement in the DMD configuration file to specify the administrator's IP address. See z/OS Communications Server: IP Configuration Reference for more information about the Exclude keyword in the DMD configuration file.

Perform the following steps if a stack's defensive filters are not blocking traffic:

6. Ensure that the filter's mode is set to Block. The Action field on the **ipsec -F display** report should indicate Defensive Block. If the Action field indicates Defensive Simulate, issue **ipsec -F update** to change the filter's mode. See z/OS Communications Server: IP System Administrator's Commands for more information about the **ipsec -F** command.
7. Ensure that the stack's mode in the DMD configuration file is set to Active. A mode of Simulate will override the individual filter's setting. It will allow a packet to match a defensive filter, generate a message, and then continue to be processed. The mode must be set to Active to cause the individual filter's mode setting to be honored. See z/OS Communications Server: IP Configuration Reference for more information about specifying a stack and its mode in the DMD configuration file. The MODIFY DISPLAY command can be issued for the DMD to display the active configuration settings. See z/OS Communications Server: IP System Administrator's Commands for more information about the MODIFY command

Perform the following step if a defensive filter is discarding traffic that should be permitted:

8. Delete the defensive filter if it is causing legitimate traffic to be discarded. Issue the **ipsec -F delete** command to delete the defensive filter from the stack. See z/OS Communications Server: IP System Administrator's Commands for more information about the **ipsec -F** command.

Perform the following steps to disable defensive filtering for a stack:

9. Update the DMD configuration file to disable defensive filters for a stack. Specify a mode of Inactive for the stack on the DmStackConfig statement. See z/OS Communications Server: IP Configuration Reference for more information about the DMD configuration file.
10. Issue the MODIFY REFRESH command for the DMD. See z/OS Communications Server: IP System Administrator's Commands for more information about the MODIFY command.

    Tip: If you are unable to update your DMD configuration file, the MODIFY FORCE_INACTIVE command can be issued for the DMD to disable defensive filtering for the stack. A later MODIFY REFRESH will use the DMD configuration file. If you want defensive filtering to remain disabled, you should update the DMD configuration file as soon as possible.

    Tip: Removing the DmStackConfig statement from the DMD configuration file will not delete existing defensive filters from the stack. If you removed the DmStackConfig statement, the defensive filters will remain in the stack until expiration. To remove the defensive filters from the stack immediately, add the DmStackConfig statement back to the DMD configuration file with a mode of Inactive or issue the MODIFY FORCE_INACTIVE command for the stack.

Perform the following step to remove all defensive filters from a stack:

11. Perform one of the following actions:

- Issue the **ipsec -F delete -N all -p stackname** command to delete all existing defensive filters from the stack. This will also delete them from the DMD's persistent storage so the filters will not be reinstalled if the stack were to be stopped and restarted. Defensive filtering remains enabled for the stack and new filters can be added to the stack. See z/OS Communications Server: IP System Administrator's Commands for more information about the **ipsec -F** command.
- Disable defensive filtering for a stack as described earlier. This will remove all existing defensive filters from the stack and the DMD's persistent storage. It will also prevent new defensive filters from being installed in the stack.

Tip: The following actions will not remove defensive filters from the stack.

- Stopping and restarting the stack. The DMD will reinstall defensive filters when the stack is restarted.
- Stopping DMD. Existing defensive filters remain installed in the stack until expiration.

Steps for diagnosing the cause for missing IP security or defensive filter syslogd output

Determine the cause for missing IP security or defensive filter syslogd output.

Procedure

Perform the following steps:

1. Ensure that Policy Agent is running on this system if IP security policy is defined.
2. Ensure that TRMD is running for this stack on this system. Consider using TCPIP PROFILE Autolog for TRMD. See "Diagnosing TRMD problems" on page 704 for more information.
3. Ensure that syslogd is running on this system.
4. Ensure that syslogd is configured for IP security and defensive filter output. TRMD always writes IP security and defensive filter log records to the syslog local4 facility.

Table 66. IPSec messages logged by TRMD

Message	Priority
EZD0827I Remote port translated	Debug
EZD0811I Decapsulation failed (reason codes 8 and 9)	Debug
All other IPSec messages logged by TRMD.	Info

Note:

a. If IP security policy is configured to log permits and denies, TRMD sends those messages to syslogd using facility local4.

b. If IKED is configured for logging, IKED messages are sent to syslogd using facility local4 and varied priorities.

Table 67. Defensive filter messages logged by TRMD

Message	Priority
Defensive filter messages logged by TRMD	Info

Note:

a. If a defensive filter indicates that a filter match should be logged, TRMD sends those messages to syslogd using facility local4.

b. If the DMD is configured for logging, DMD messages are sent to syslogd using facility local4 and varied priorities.

Tips:

- If TRMD is logging Intrusion Detection Services (IDS) messages, as well as IP Security (IPSec) messages and defensive filter messages, consider using the facility to separate the IDS messages from the IPSec and defensive filter messages. IDS messages are written to the daemon facility. IPSec and defensive filter messages are written to the Local4 facility.
- If running multiple TRMDs, consider using the syslogd -u option when starting syslogd. The -u option causes the job name of the application writing the syslogd record to be included in the syslogd record.
- If running multiple TRMDs, consider using the trmd jobname prefix to separate IPSec output by stack.

Guidelines:

- Ensure that syslogd is configured to write TRMD and IKED messages for IP security.
- Ensure that syslogd is configured to write TRMD and DMD messages for defensive filters.
- For example, the following lines could be added to the syslogd configuration file to organize TRMD, IKED, and DMD messages:

```
*.*.local4.* /tmp/logs/filter.log
*.IKED*.local4.* /tmp/logs/IKED.log
*.DMD*.local4.* /tmp/logs/DMD.log
*.trmd*.local4.* /tmp/logs/trmdfilt.log
*.trmd*.daemon.* /tmp/logs/ids.log
```

  In the example, IKED, DMD, and TRMD IP security and defensive filter messages are all written to the log file /tmp/logs/filter.log. IKED messages are also written to the log file /tmp/logs/IKED.log. DMD messages are also written to the log file /tmp/logs/DMD.log. IP security and defensive filter TRMD messages are also written to the log file /tmp/logs/trmdfilt.log. If TRMD is logging IDS messages, those messages are written to /tmp/logs/ids.log.
- Ensure that the log files exist or syslogd is configured to create them using the -c option.
- Ensure that the log files are writable.
- Ensure that there is adequate space on the file system for writing to the log files.

  Perform the following steps to reduce the amount of syslogd output for IP security and defensive filters.

  a. Ensure that the logging levels for the IKE daemon are set appropriately in the IKE daemon configuration file.

- IkeSyslogLevel - During day-to-day operation, this value should be set no higher than the default of 1. A higher value should be used for temporary diagnostic purposes only. IkeSyslogLevel can also be set to 0 to disable IKE syslog messages entirely.
- PagentSyslogLevel - During day-to-day operation, this value should be set to the default of 0. A higher value should be used for temporary diagnostic purposes only.

b. Ensure that the logging levels for the DMD are set appropriately in the DMD configuration file.
   - SyslogLevel - During day-to-day operation, this value should be set no higher than 7. A higher value should be used for temporary diagnostic purposes only. SyslogLevel can also be set to 1 for minimum logging or 0 to disable DMD syslog messages entirely.

c. Ensure that filter logging controls are set appropriately for IP security filters.
   - Filter logging generates a message each time an inbound or outbound packet matches the filter. Exhaustive logging of IP traffic can have a negative effect on performance. Filter logging can be controlled at the individual rule level, including the ability to specify whether to log permitted traffic, denied traffic, or both.
   - To disable filter logging for profile filter rules:
     - To disable logging for a configured filter rule, set NOLOG on the IPSECRULE or IPSEC6RULE statement.
     - To disable logging for the implicit filter rules that deny all traffic not permitted by a configured rule, set NOLOGIMPLICIT on the IPSEC statement.
     - To disable filter logging for all profile filter rules, set LOGDISABLE on the IPSEC statement.
   - To disable filter logging for policy filter rules configured using the Policy Agent:
     - To disable logging for a configured filter rule, set IpFilterLogging No on the IpGenericFilterAction statement.
     - To disable logging for the implicit filter rules that deny all traffic that does not match a configured rule, set IpFilterLogImplicit No on the IpFilterPolicy statement.
     - To disable filter logging for all policy filter rules, set FilterLogging Off on the IpFilterPolicy statement.
   - To disable filter logging for policy filter rules configured with the IBM Configuration Assistant for z/OS Communications Server:
     - To disable logging for a configured filter rule, set filter logging to No for the Connectivity Rule.
     - To disable logging for the implicit filter rules that deny all traffic that does not match a configured rule, select Do NOT log implicit deny events on the IPSec: Stack Level Settings panel.
     - To disable filter logging for all policy filter rules, select **Disable all filter logging** on the IPSec: Stack Level Settings panel.
   - The following messages are controlled by the configured filter logging settings described above:
     - EZD0814I Packet permitted
     - EZD0815I Packet denied by policy
     - EZD0821I Packet denied, no tunnel

- EZD0822I Packet denied, tunnel inactive
- EZD0832I Packet denied by NAT Traversal Processing
- EZD0833I Packet denied, tunnel mismatch

d. Ensure that filter logging controls are set appropriately for defensive filters.
   - Filter logging generates a message each time an inbound or outbound packet matches the defensive filter. Exhaustive logging of IP traffic can have a negative effect on performance. Filter logging can be controlled at the individual defensive filter rule level.
   - To disable filter logging for a defensive filter rule, use the **ipsec -F update** command with log no specified.
   - To limit filter logging for a defensive filter rule, use the **ipsec -F update** command with the loglimit keyword specified with a value of 1 - 9999. See the **ipsec -F** command in z/OS Communications Server: IP System Administrator's Commands for more information about the **ipsec -F update** command.
   - The following messages are controlled by the defensive filter's log and loglimit settings:
     - EZD1721I Packet denied by defensive filter
     - EZD1722I Packet would have been denied by defensive filter

e. Ensure that IP security and defensive filter messages being logged by the TRMD daemon are being handled appropriately.
   - The TCP/IP stack invokes the TRMD daemon to log IP security and defensive filter messages to syslog. The filter logging messages described above are logged by the TRMD daemon. TRMD also logs messages that are not associated with a specific filter. For example, when a tunnel is successfully negotiated, TRMD logs message "EZD0818I Tunnel added". Also, when an IP security policy update is processed, TRMD logs message "EZD0816I IPSec Policy updated". When a defensive filter is added to the stack, TRMD logs message "EZD1723I Defensive filter added".
   - There is no explicit configuration option to turn off logging for TRMD messages that are not associated with a specific filter. However, the syslog configuration file can be updated to exclude some or all TRMD messages. See Table 66 on page 737 for information about the syslog priority used to log TRMD messages. See Table 67 on page 738 for information about the syslog priority used to log defensive filter TRMD messages.
   - Include the following line in your syslog configuration file to exclude IP security TRMD messages logged with a priority of debug. IP security and defensive filter TRMD messages with a priority of info or higher would be written to /tmp/trmdlog. Messages with a priority of debug would not be written to the file.

     ```
     *.TRMD*.local4.info  /tmp/trmdlog
     ```

   - Include the following line in your syslog configuration to exclude all IP security and defensive filter TRMD messages.

     ```
     *.TRMD*.*.*;*.TRMD*.local4.none   /tmp/trmdlog
     ```

     All messages with job name TRMD* would be selected. Then all TRMD messages using facility local4 would be excluded. In effect this excludes all IP security and defensive filter TRMD messages from being written to /tmp/trmdlog.

Steps for verifying IP security and defensive filter operation

Verify IP security and defensive filters operations.

Before you begin

Identify the characteristics of the IP traffic for which IP security operation is to be verified. The characteristics of IP traffic that are subject to IP security control are described by the IpFilterRule or IPSECRULE (for IPv4) or IPSEC6RULE (for IPv6) statement. See z/OS Communications Server: IP Configuration Reference for more information about the IpFilterRule, IPSECRULE and IPSEC6RULE statements.

About this task

Figure 94 on page 742 shows the decisions involved for IP security operation.

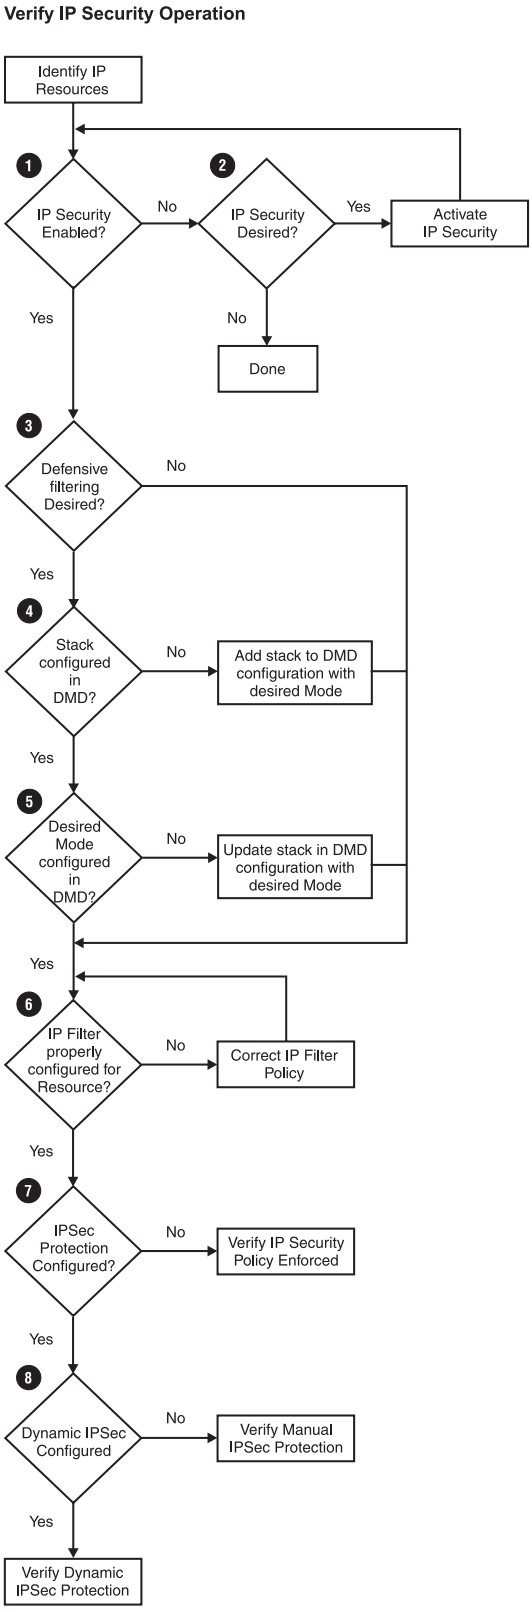

Figure 94. Overview of verifying IP security operation

Procedure

Perform the following steps:

1. Use the Netstat CONFIG/-f command to determine whether the TCP/IP stack is configured for IP security for IPv4, IPv6, or both. For information about the Netstat command, see z/OS Communications Server: IP System Administrator's Commands.

   Do one of the following:

   - If the stack is not configured for IP security for the IP protocol that you want, proceed to step 2.
   - If the stack is configured for IP security for the IP protocol that you want, proceed to step 3.

2. If you want IP security enabled for IPv4, configure the stack for IPv4 IP security using the IPCONFIG IPSECURITY statement in the TCP/IP profile. If you want IP security enabled for IPv6, configure the stack for IPv6 IP security using the IPCONFIG6 IPSECURITY statement in the TCP/IP profile. See z/OS Communications Server: IP Configuration Reference for more information about the IPCONFIG IPSECURITY and IPCONFIG6 IPSECURITY statements. See z/OS Communications Server: IP Configuration Guide for general information about IP security concepts, including IP filtering.
3. Use the MODIFY command to display the configuration values for the Defense Manager daemon (DMD). For more information about the MODIFY command, see z/OS Communications Server: IP System Administrator's Commands.

   If you want defensive filtering enabled, do one of the following:

   - If the stack name is not listed in the DMD configuration , proceed to step 4.
   - If the stack name is listed in the DMD configuration but does not have the mode that you want, proceed to step 5.
   - If the stack name is listed in the DMD configuration with the mode that you want, proceed to 6.

   Otherwise, if you do not want defensive filtering, proceed to 6.
4. Update the DMD configuration file to include a DmStackConfig statement for the stack for which you want defensive filtering. Specify the defensive filtering mode, Active or Simulate, on the DmStackConfig statement with the stack name. See z/OS Communications Server: IP Configuration Reference for more information about the DMD configuration file. Proceed to 6.
5. Update the DMD configuration file to specify the defensive filtering mode, Active or Simulate, on the DmStackConfig statement. See z/OS Communications Server: IP Configuration Reference for more information about the DMD configuration file.
6. Use the **ipsec -t** command to determine which IP filter applies to the identified IP packet. At the top of the **ipsec -t** command output, note whether Source indicates Stack Profile or Stack Policy.

   Limited IP filter controls can be configured using the IPSECRULE statement (for IPv4) and the IPSEC6RULE statement (for IPv6) in the TCP/IP profile. Full IP security capability, including manual and dynamic IPSec protection, requires use of the Policy Agent for IP security policy configuration.

   Locate the Type field in the **ipsec -t** command output to determine the type of filter. If the Type field indicates Defensive, then the filter is a defensive filter. Defensive filters are not configured but are added to the stack by the **ipsec** command. Typically, this is done by an external security information and event manager that detects an attack. However, the **ipsec** command can be issued manually by a user with the appropriate authority to add a defensive filter.

Tip: The **ipsec -t** command can return multiple filter rules because the actual packet filtering compares more attributes than might be supplied as input on the **ipsec -t** command. To minimize this effect, supply as much information as possible on the **ipsec -t** command.

If the returned filter rules include a defensive filter, take the following actions:

- Locate the exclusion list at the top of the **ipsec -t** command output and determine whether there are any IP addresses listed. Traffic from IP addresses in the exclusion list will bypass defensive filters.
- Locate the Action field in the **ipsec -t** command output to determine the mode of the defensive filter. If the Action field indicates `Defensive Block` the filter is discarding traffic. If the Action field indicates `Defensive Simulate` only filter logging is done, packets continue to be processed.
- If the defensive filter rule is blocking traffic that should be allowed, determine the user that added the filter by inspecting the syslog messages. Locate the "EZD1723I Defensive filter added" defensive filter message that corresponds to this defensive filter. The userid of the user that added the filter is included in the message.

If none of the filters that are returned by the **ipsec -t** command include the action for the identified IP packet that you want, then correct the IP filter configuration. See z/OS Communications Server: IP Configuration Guide for general information about configuring IP filters.

7. Locate the Type field in the **ipsec -t** command output to determine whether IPSec protection is configured for the identified IP packet. If the Type field indicates `Generic` or `Defensive`, then IPSec protection is not configured for the identified IP packet. See "Steps for verifying IP security policy or defensive filter enforcement" on page 749 to verify that the configured policy is enforced for the IP traffic characterized by the identified IP packet.
8. Locate the Type field in the **ipsec -t** command output to determine whether manual or dynamic IPSec protection is configured for the identified IP packet. If the Type field indicates Manual, then see "Steps for verifying manual IPSec protection." If the Type field indicates Dynamic or Dynamic Anchor, then see "Steps for verifying dynamic IPSec protection" on page 746.

Results

Steps for verifying manual IPSec protection

Verify manual IPSec protection.

Before you begin

Complete the steps in "Steps for verifying IP security and defensive filter operation" on page 741 in order to identify the name of an IpFilterRule for which manual IPSec protection is to be verified.

About this task

Figure 95 on page 745 shows the decisions involved for verifying manual IPSec protection.

Verify Manual IPSec Protection

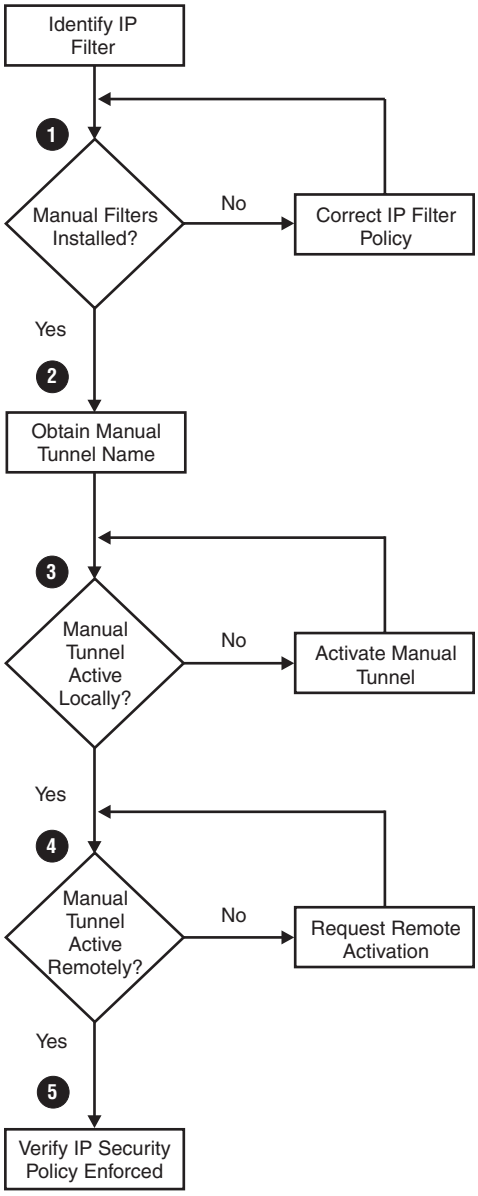

Figure 95. Overview of verifying manual IPSec protection

Procedure

Perform the following steps:

1. Verify that manual filters that correspond to the identified IpFilterRule are installed in the stack by using the **ipsec -f display -n** command. Two filters of type Manual (1 inbound and 1 outbound) are installed in the stack for an IpFilterRule that is configured with IpManVpnAction. If the manual filter rules are not installed in the stack, then correct the IP filter policy. An IpFilterRule might be inactive (not installed) in the stack due to an IpTimeCondition. For information about the **ipsec** command, see z/OS Communications Server: IP System Administrator's Commands. See z/OS Communications Server: IP Configuration Reference for more information about the IpManVpnAction and IpTimeCondition statements.

If IP filter rules are not installed, also verify that Policy Agent is active.

2. Obtain the IpManVpnAction name by locating the VpnActionName field in the **ipsec -f** command output. This is the name of the IpManVpnAction policy configuration statement. Obtain the manual tunnel ID by locating the TunnelID field in the **ipsec -f display** command output. The Tunnel ID for a manual tunnel has a value of `M`, followed by a positive integer.
3. Verify that the manual tunnel is active.

   Use the **ipsec -m display -a** command, supplying the manual tunnel ID.

   Locate the State field in the **ipsec -m** command output and confirm that it indicates Active. If the manual tunnel is not active, then activate the tunnel using the **ipsec -m activate** command. You might consider updating the IpManVpnAction policy configuration statement to specify **Active yes**, if it is not already specified. A setting of **Active yes** causes the manual tunnel state to be set to active when the manual tunnel is installed in the stack, without the additional step of issuing ipsec -m activate.

   If you are using the IBM Configuration Assistant for z/OS Communications Server to configure, you can choose to automatically activate manual tunnels within each Connectivity Rule.
4. Contact the remote security endpoint's network administrator to ensure that the manual tunnel has been activated remotely. For traffic to flow through a manual tunnel, the remote security endpoint must also activate the manual tunnel.
5. Verify that IpManVpnAction is enforced. See "Steps for verifying IP security policy or defensive filter enforcement" on page 749.

Steps for verifying dynamic IPSec protection

Verify dynamic IPSec protection.

Before you begin

Complete the steps in "Steps for verifying IP security and defensive filter operation" on page 741 in order to identify the name of an IpFilterRule for which dynamic IPSec protection is to be verified.

About this task

Figure 96 on page 747 shows the decisions involved for verifying dynamic IPSec protection.

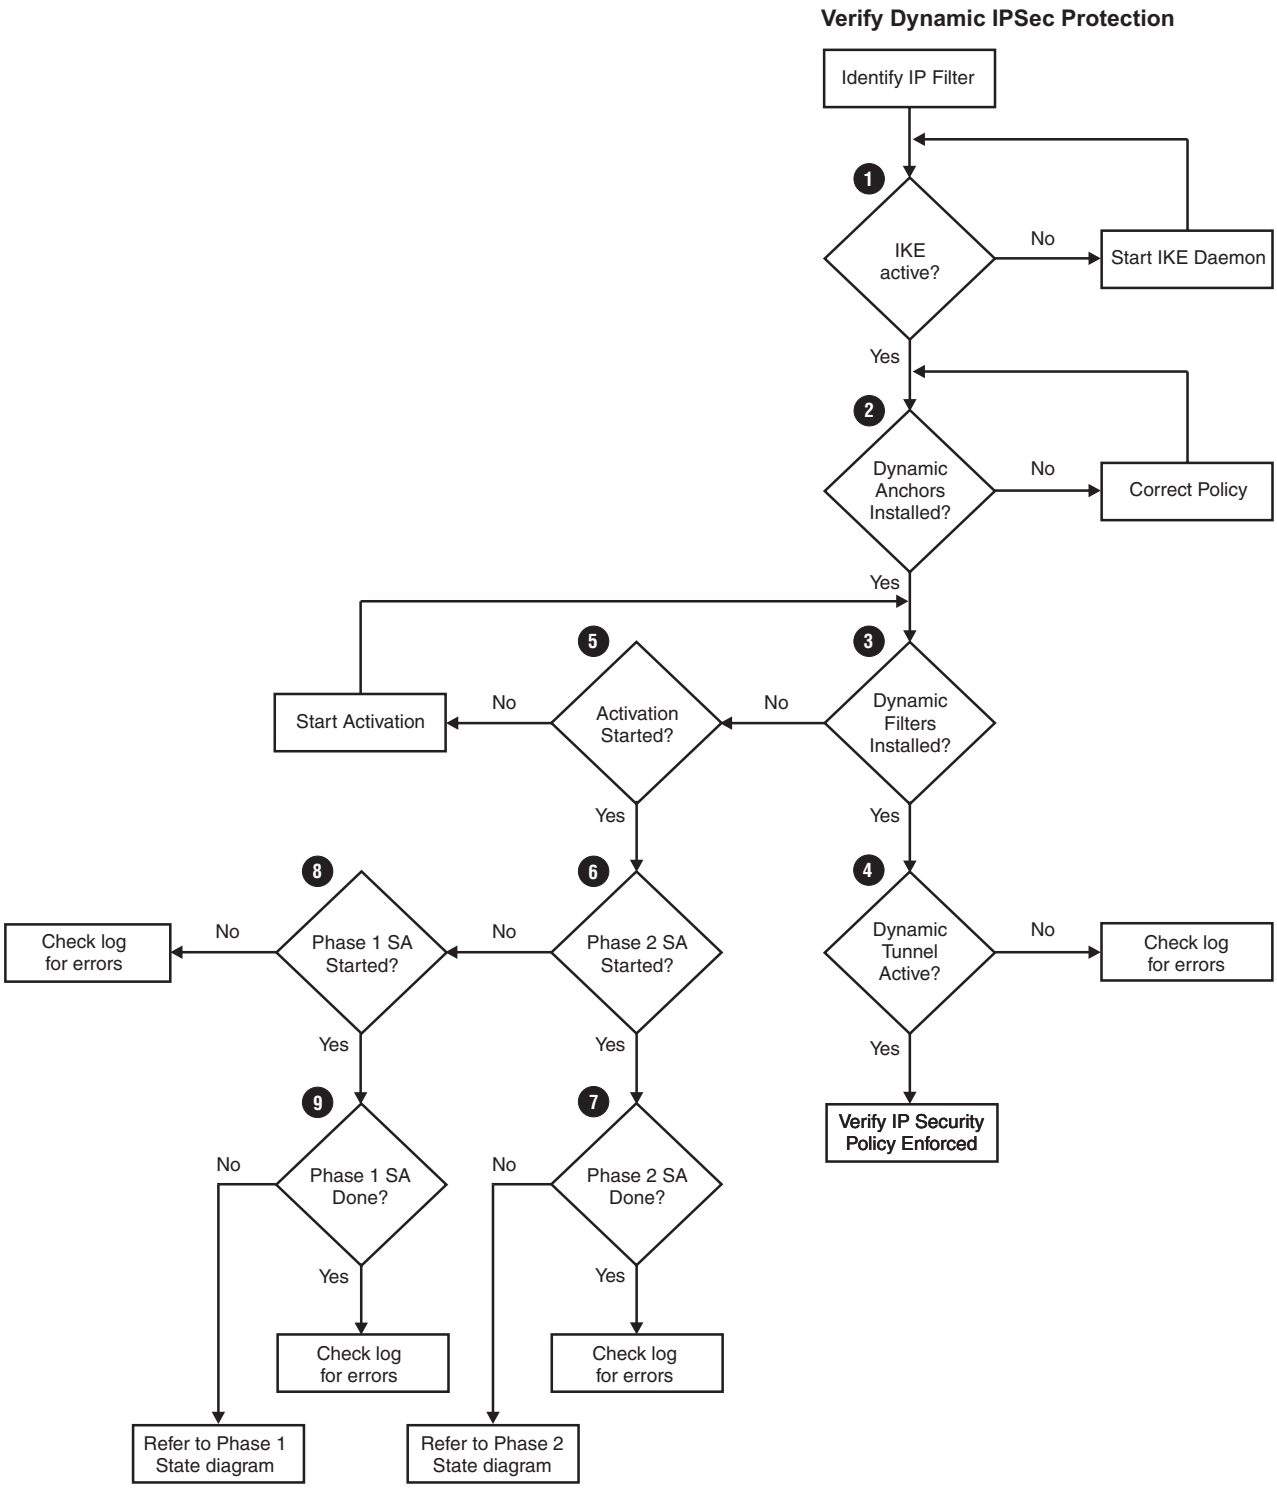

Figure 96. Overview of verifying dynamic IPSec protection

Procedure

Perform the following steps:

1. Verify that the IKE daemon is active. See "Steps for verifying server operation" on page 27.

Tip: The IKE daemon binds to UDP ports 500 and 4500.

2. Use the **ipsec -f display -n** command to verify that dynamic anchor filters that corresponds to the identified IpFilterRule are installed in the stack. Two filters of type Dynamic Anchor (1 inbound and 1 outbound) are installed in the stack for an IpFilterRule that is configured with an IpDynVpnAction. If the dynamic anchor filter rules are not installed in the stack, then correct the IP filter policy. An IpFilterRule might be inactive (not installed) in the stack due to an IpTimeCondition. For information about the **ipsec** command, see z/OS Communications Server: IP System Administrator's Commands. See z/OS Communications Server: IP Configuration Reference for more information about the IpDynVpnAction and IpTimeCondition statements. If IP filter rules are not installed, also check the following:
   - Verify that policy agent is active.
   - If policy agent is active, verify that the following messages appeared after IKED was started:

     ```
     EZD1058I IKE STATUS FOR STACK stackname IS UP
     EZD1068I IKE POLICY UPDATED FOR STACK stackname
     ```

   If these messages did not appear, check the Policy Agent log for errors.
3. Use the **ipsec -f display -n** command to verify that the dynamic filters are installed in the stack. When the IKE daemon completes a dynamic tunnel negotiation, it installs two dynamic filters to more specifically control the IP traffic that can be permitted through the dynamic tunnel.

   The dynamic filters are identified with a Type field of Dynamic in the **ipsec** command output.

   Do one of the following:
   - If no dynamic filters are installed in the stack with the identified IpFilterRule name, then proceed to step 5.
   - If the dynamic filters are installed in the stack, then proceed to step 4.
4. Verify that the dynamic tunnel that corresponds to the dynamic filters is active.

   The IKE daemon installs a dynamic tunnel and corresponding inbound and outbound dynamic filters into the stack.

   Follow these steps to perform verification:
   a. Locate the dynamic tunnel ID in the TunnelID field of the **ipsec -f** command output.

      Tip: Be sure to look for the TunnelID identified on the filter rule with type Dynamic, rather than the filter rule with type Dynamic Anchor.
   b. Use the **ipsec -y display -a** command, supplying the dynamic tunnel ID.
   c. Locate the State field in the **ipsec -y** command output and confirm that it indicates Active. If the dynamic tunnel is not active, then check the IKE syslogd output for errors. Otherwise, see "Steps for verifying IP security policy or defensive filter enforcement" on page 749.
5. If no dynamic filters have been installed in the stack, then the dynamic tunnel activation might not have been started.

   Consider whether you need to activate the tunnel.
   - If you intend to manually start the tunnel, then you must issue the **ipsec -y activate** command. If you intend for the tunnel to be automatically activated, you must configure your LocalDynVpnPolicy to include a LocalDynVpnRule with AutoActivate **Yes** specified.
   - If you intend for the tunnel to be activated on demand by outbound traffic, then you must configure AllowOndemand **Yes** on either your IpFilterPolicy

or on an IpLocalStartAction associated with the IpFilterRule identified in step 2 on page 748, and you must also set the outbound traffic flow to trigger the activation.

- If the tunnel is intended to be activated by the remote security endpoint, then you must configure your KeyExchangePolicy properly, and the remote security endpoint must initiate the tunnel negotiation. If you know that you have not yet taken a required action to activate the tunnel, do so now. Otherwise, proceed to the next step.

See z/OS Communications Server: IP Configuration Guide for more information about activating dynamic tunnels.

6. Use the **ipsec -y display -b** command to display all dynamic tunnels known to the IKE daemon. In the **ipsec** command output, search for a dynamic tunnel with an IpFilterRule name that matches the identified IpFilterRule name. If there is no such dynamic tunnel, proceed to step 8. Otherwise, proceed to step 7.
7. If the state of the dynamic tunnel that was identified in step 6 is not DONE, then see "Interpreting IKEv1 daemon phase 2 SA states" on page 980 or "Interpreting IKEv2 Child SA states" on page 989. Otherwise, check the syslogd output for errors.
8. Use the **ipsec -k display** command to see whether there is an applicable IKE tunnel negotiation in progress.

   If not, check the log for errors. Otherwise, proceed to step 9.
9. If the IKE tunnel state is not DONE, then note the role (initiator or responder) of the IKE tunnel and see "Interpreting IKEv1 daemon phase 1 SA states" on page 974 or "Interpreting IKEv2 IKE SA states" on page 986. Otherwise, check the syslogd output for errors.

Steps for verifying IP security policy or defensive filter enforcement

Verify IP security policy or defensive filter enforcement.

Before you begin

Complete the steps in "Steps for verifying IP security and defensive filter operation" on page 741 in order to identify the name of an IpFilterRule or IPSECRULE or IPSEC6RULE for which IP security policy enforcement is to be verified or the name of a defensive filter.

About this task

Figure 97 on page 750 shows the decisions involved for verifying IP security policy enforcement.

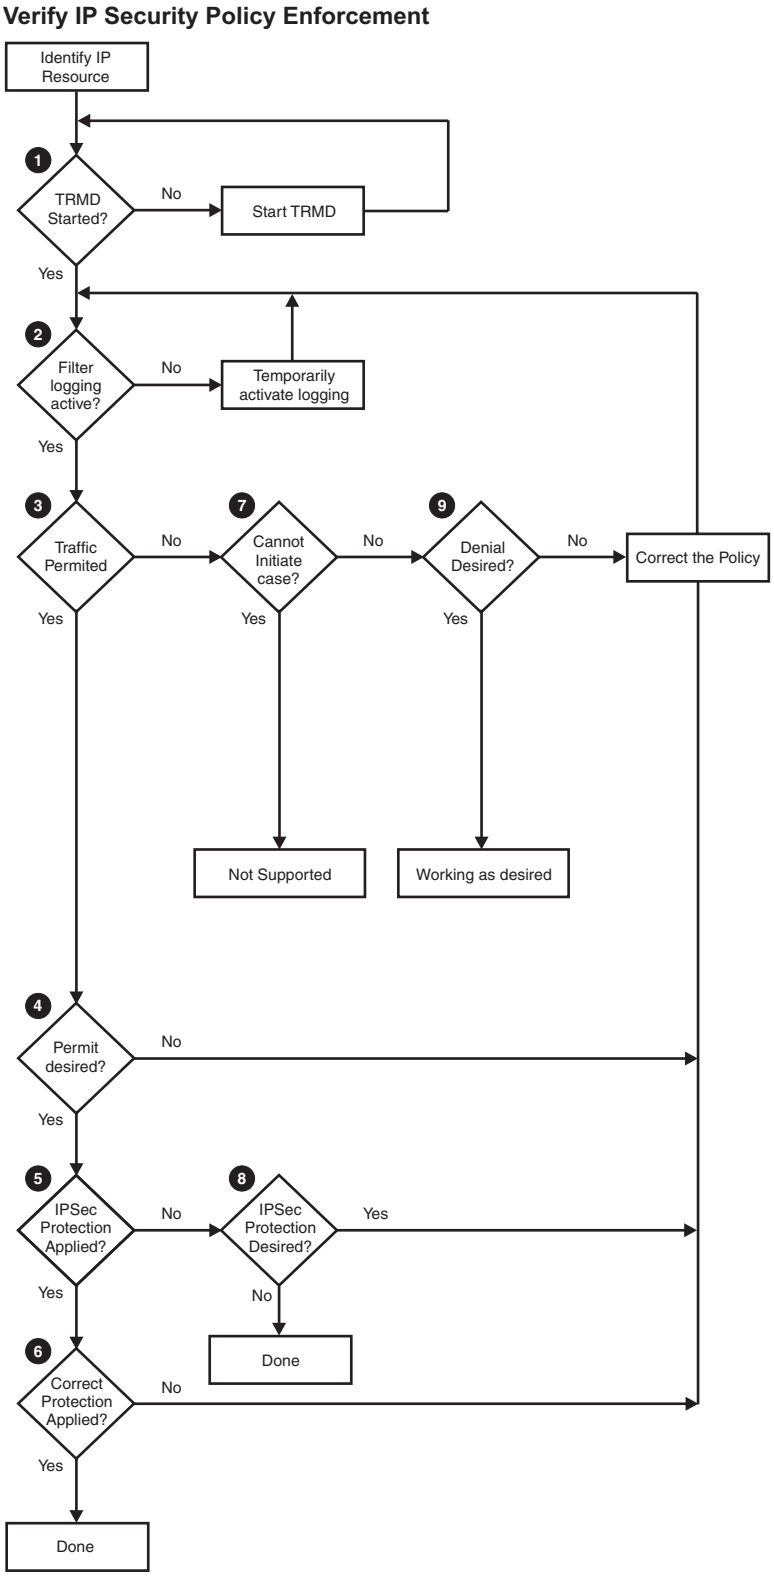

Figure 97. Overview of verifying IP security policy enforcement

Procedure

Perform the following steps:

1. Start TRMD for the stack if it is not already active. The Traffic Regulation Management Daemon (TRMD) is required to log IP filter permits and denies. See z/OS Communications Server: IP Configuration Reference for information about starting TRMD.
2. Display the identified filter rule using the **ipsec -f display -n** command if it is an IP security filter. Display the identified filter rule using the **ipsec -f display -N** command if it is a defensive filter. Use the instructions in the following lists to temporarily activate logging for the filter if it is not already active. See z/OS Communications Server: IP Configuration Reference for information about the IpFilterPolicy, IpFilterRule, IpGenericFilterAction, IPSEC, IPSECRULE, and IPSEC6RULE statements.
   - If the displayed filter Type field is `Defensive`, do the following:
     - If the displayed filter Logging field is not `ALL`, update the defensive filter's log setting with the **ipsec -F update** command. See z/OS Communications Server: IP System Administrator's Commands for information about the **ipsec** command.
   - If the **ipsec** command header output indicates Stack Policy, do the following:
     - If the **ipsec** command header output indicates `Logging NO`, temporarily specify `FilterLogging On` on the IpFilterPolicy statement. If you are using the IBM Configuration Assistant for z/OS Communications Server to configure, select **Enable filter logging** on the IPSec: Stack Level Settings panel.
     - If the displayed filter Logging field is not `ALL`, specify IpFilterLogging **Yes** on the IpGenericFilterAction referenced by the IpFilterRule. If you are using the IBM Configuration Assistant for z/OS Communications Server to configure, set filter logging to **Yes** in each Connectivity Rule.
     - Use the MODIFY command with the Policy Agent to activate your changes, if any. See z/OS Communications Server: IP System Administrator's Commands for more detailed information about the MODIFY command.
   - If the **ipsec** command header output indicates Stack Profile, do the following:
     - If the **ipsec** command header output indicates Logging NO, specify LOGENABLE on the IPSEC statement.
     - If the displayed filter Logging field does not indicate ALL, specify LOG on the IPSECRULE or IPSEC6RULE statement.
     - Use the VARY TCPIP,,OBEYFILE command to activate your changes, if any.
3. After IP filter logging is active, check the syslog to determine whether the IP traffic that is characterized by the filter rule is being permitted or denied. Message EZD0814I is issued when an IP packet is permitted. Message EZD0815I, EZD0821I, EZD0832I, EZD0822I, EZD0833I, or EZD1721I is issued when an IP packet is denied. If the traffic is denied, proceed to step 7 on page 752. Otherwise, proceed to step 4.
4. If the IP traffic is being permitted, but you do not want that, correct the filter configuration. See z/OS Communications Server: IP Configuration Guide for information about configuring IP filtering.
5. Determine whether the IP traffic is subject to IPSec protection by locating the vpnaction field in the EZD0814I message. If the vpnaction field is not N/A then the IP traffic is subject to IPSec protection. If IPSec protection is not applied, then proceed to step 8 on page 752. Otherwise, proceed to step 6 on page 752.

6. Determine the properties of the IPSec tunnel by first locating the tunnelID field in the EZD0814I message. Apply the following criteria to evaluate the tunnelID:
   - If the first character of the tunnelID is M, use the **ipsec -m display -a** command to display the corresponding manual tunnel. If the displayed manual tunnel does not have the characteristics that you want, correct the IpManVpnAction statement. See z/OS Communications Server: IP Configuration Reference for information about the IpManVpnAction statement.

     If you are using the IBM Configuration Assistant for z/OS Communications Server to configure, the IpManVpnAction corresponds to a Security Level implementing Manual Tunnels. If Security Level does not contain the characteristics that you want, correct the Security Level. See the Configuration Assistant online help for additional information.
   - If the first character of the tunnelID is Y, use the **ipsec -y display -a** command to display the corresponding dynamic tunnel. If the displayed dynamic tunnel does not have the characteristics that you want, correct the IpDynVpnAction statement. See z/OS Communications Server: IP Configuration Reference for information about the IpDynVpnAction statement.

     If you are using the IBM Configuration Assistant for z/OS Communications Server to configure, the IpDynVpnAction corresponds to a Security Level implementing dynamic tunnels. If Security Level does not contain the characteristics that you want, correct the Security Level. See the Configuration Assistant online help for additional information.
7. Data traffic cannot be initiated to the remote data endpoint in certain cases when NAT traversal support is being used. Message EZD0832I is issued when an attempt is made to initiate data traffic if either of these conditions is true:
   - The remote security endpoint is acting as a security gateway and a NAT was detected between the local security endpoint and the remote security endpoint.
   - The remote security endpoint is behind a NAT device performing port translation

   If not the "cannot initiate case" message, proceed to 9.
8. If the IP traffic is not being protected with IPSec, but you want IPSec protection, correct the filter configuration. See z/OS Communications Server: IP Configuration Guide for information about configuring IP filtering.

   Tips:
   - If you defined a filter rule for this traffic but packets are matching a different filter rule, consider whether your filter rules are ordered correctly. The filter table is searched in sequential order, so it is possible that a filter rule earlier in the table is overshadowing a filter rule later in the table.
   - If you defined a filter rule for this traffic and packets are matching your filter rule, consider whether you defined the right policy action for the filter rule. You can choose an action of deny, permit, or IPSec protection, with further options for IPSec protection.
9. If the IP traffic is being denied, determine what type of filter is denying the traffic.

   If the IP traffic is being denied by a defensive filter, message EZD1721I is issued.

- Determine the user that added the defensive filter by locating the "EZD1723I Defensive filter added" message that corresponds to the deny message. The userid of the user that added the filter is included in message EZD1723I.
- Delete the defensive filter if it is denying traffic that should be permitted. Use the **ipsec -F delete** command to delete a defensive filter. See z/OS Communications Server: IP System Administrator's Commands for information about the **ipsec** command.

If the IP traffic is being denied by an IP security filter, correct the filter configuration to change this situation. See z/OS Communications Server: IP Configuration Guide for information about configuring IP filtering.

Tips:

- If you defined a filter rule for this traffic but packets are matching a different filter rule, consider whether your filter rules are ordered correctly. The filter table is searched in sequential order, so it is possible that a filter rule earlier in the table is overshadowing a filter rule later in the table.
- If you change the order of two filter rules that employ dynamic IPSec protection, consider whether a similar change is necessary in your key exchange policy. For more information about the ordering of key exchange policy rules, see z/OS Communications Server: IP Configuration Guide.
- If you defined a filter rule for this traffic and packets are matching your filter rule, consider whether you defined the right policy action for the filter rule. You can choose an action of deny, permit, or IPSec protection, with further options for IPSec protection.

Steps for verifying IPSec processing on zIIP

If attempting to use the zIIP IPSECURITY support (to direct IPSec AH|ESP protocol processing to zIIP), issue the Netstat STATS/-S command while the IPSec workload is running. The inbound and outbound 'Packets Handled by zIIP' counters will be rising if IPSec workload is in fact being processed on zIIP(s). If these counters are not rising while IPSec traffic is flowing, verify (a) GLOBALCONFIG ZIIP IPSECURITY parameters are specified in the TCPIP profile (use Netstat Config/-f to verify); and (b) zIIP(s) are configured to the z/OS image (use MVS D M=CPU command to verify).

Determining the Workload Manager service class associated with IPSec workload being processed on zIIP

To verify that the new independent enclave is being used with an appropriate WLM service class issue the SDSF ENC command or view the RMF Workload Activity report. For more information regarding the SDSF function of viewing enclaves, see z/OS SDSF Operation and Customization. For additional information regarding the RMF Workload Activity report, see z/OS RMF Report Analysis.

Tools for diagnosing IP security and defensive filter problems

This topic describes tools used to diagnose IP security and defensive filter problems.

Using the ipsec command

You can use the **ipsec** command to display information about:

- IP filter rules
- Security associations

- Port translation
- SECCLASS definitions
- Matching IP filter rules for a specified traffic pattern
- Network security information of an IKE daemon's active NSS IPSec clients
- NSS IPSec clients connected to NSS servers

By default, **ipsec** commands are directed to the local system. Optionally, **ipsec** commands may be directed to remote systems (NSS IPSec clients) using the **ipsec -z** option.

Restriction: Management of defensive filters (**ipsec -F**) is only provided through the local **ipsec** command. Remote management using an NSS server is not supported.

ipsec -f display

The **ipsec -f** display command displays information about the current set of filter rules in use by a stack. The current set of filter rules will include any installed defensive filters.

You can use the options listed in Table 68 to define the display.

Table 68. **ipsec -f** *display command options*

Option	Use
-p	Directs the command to a stack other than the local default stack.
-z	Directs the command to an NSS IPSec client.
-c profile	Displays information about the set of filter rules defined on the IPSEC statement in the TCP/IP profile.
-c policy	Display information about the set of filter rules defined in the Policy Agent IPSec Configuration file.

Filter rules that are disallowed due to time conditions do not appear in the output of **ipsec -f** display command. The **pasearch** command must be used to obtain information about such filter rules. When working with an NSS client the **pasearch** command needs to be issued on the system where the client is executing. Use the **ipsec -x** command to determine where the NSS IPSec client is executing.

Several different types of filter rules exist. By default, the **ipsec -f** display output includes information about generic, defensive, dynamic anchor, dynamic, NATT anchor, and NATT dynamic filter rules. You can use the -h option to display information about filter rules of type NRF. NAT resolution filter (NRF) rules are present when the remote security endpoint is behind a NAT. See z/OS Communications Server: IP Configuration Guide for an explanation of filter types.

ipsec -F display

The **ipsec -F display** command displays information about defensive filters.

You can use the options listed in Table 69 on page 755 to define the display.

Table 69. **ipsec -F** *display command options*

Option	Use
-p	Directs the command to a stack other than the local default stack.
-G	Directs the command to the Defense Manager daemon (DMD) to display global defensive filters.

Several different types of filter rules exist. The **ipsec -F display** output only includes defensive filters.

ipsec -F update and **ipsec -F delete**

The **ipsec -F update** command can be used to update a defensive filter. The **ipsec -F delete** command can be used to delete a defensive filter.

You can use the options listed in Table 70 to define the update and delete values.

Table 70. **ipsec -F update** *or* **delete** *command options*

Option	Use
-p	Directs the command to a stack other than the local default stack.
-G	Directs the command to the Defense Manager daemon (DMD) to manage global defensive filters.
-N	Provides the name of the defensive filter to be updated or deleted.

ipsec -m display

The **ipsec -m** display command displays information about manual tunnels installed in the stack. Use the -p option to direct the command to a stack other than the default stack or the -z option to direct the command to an NSS IPSec client. Manual tunnels can be either active or inactive. A manual tunnel must be active before traffic matching a filter rule that uses the manual tunnel can be permitted.

Manual tunnels that are not allowed to be used due to time conditions do not appear in the output of **ipsec -m** display command. Use the **pasearch** command to obtain information about such manual tunnels.

ipsec -k display

The **ipsec -k** display command displays information about IKE tunnels for the default stack. This information is obtained from the IKE daemon. Use the **-p** option to direct the IKE daemon to return information about a different stack or the **-z** option to direct the command to an NSS client.. An IKE tunnel must be in place before a dynamic IPSec (phase 2) security association can be negotiated by the IKE daemon.

Sometimes multiple phase 1 security associations that correspond to the same IKE tunnel can occur. By default, only information about the most current phase 1 security association for an IKE tunnel is displayed. Use the **-c** option to display information about all phase 1 security associations corresponding to an IKE tunnel.

Security associations for use by a dynamic tunnel are negotiated under the protection of a phase 1 security association. Specify the **-e** option to display information about IPSec security associations that were negotiated or that are in the process of being negotiated under the protection of a phase 1 security association.

ipsec -y display

The **ipsec -y** display command displays information about dynamic tunnels installed in the default stack. Use the **-p** option to direct the command to another stack or the **-z** option to direct the command to an NSS IPSec client. A dynamic tunnel must be active before traffic matching a filter rule utilizing an IpDynVpnAction can be permitted.

At times, there might be multiple IPSec security associations that correspond to the same dynamic tunnel. By default, only information about the most current IPSec security association for a dynamic tunnel is displayed. Use the **-c** option to display information about all IPSec security associations that correspond to a dynamic tunnel

The stack only knows about IPSec security associations that have been successfully negotiated. The IKE daemon knows about IPSec security associations that have been successfully negotiated as well as those currently being negotiated. At times, it is helpful to see information about IPSec security associations that are in the process of being negotiated. The **-b** option obtains information about IPSec security associations from the IKE daemon rather than the stack.

When a stack is a target for a distributed DVIPA it might contain IPSec security associations for a dynamic tunnel that was negotiated on behalf of the distributing stack. Such security associations are known as shadow security associations. The **-s** option obtains information about shadowed security associations.

ipsec -i

Use the **ipsec -i** command to display the SECCLASS value assigned to interfaces defined to the default stack. The **-p** option directs the command to another stack or the **-z** option directs the command to an NSS IPSec client. The SECCLASS option of the LINK or INTERFACE statement is used to assign a security classification to an interface. The LINK or INTERFACE statement is specified in the TCP/IP profile. SECCLASS can be specified as a filtering criteria on certain IP filter rules.

ipsec -t

Use the **ipsec -t** command to locate active filter rules for the default stack that match a specified traffic pattern. The **-p** option directs the command to another stack or the **-z** option directs the command to an NSS IPSec client.

ipsec -o

Use the **ipsec -o** command to display the port translation table of the default stack. The **-p** option directs the command to another stack or the **-z** option directs the command to an NSS IPSec client. Port translation is done as needed for TCP and UDP connections that use a dynamic security association with a remote security endpoint that is behind a NAT.

ipsec -w

Use the **ipsec -w** command to display network security information for each of an IKE daemon's active NSS IPSec clients.

ipsec -x

Use the **ipsec -x** command to display a list of the NSS IPSec clients connected to the NSS server. Each NSS IPSec client represents a remote system made up of an IKE daemon and a TCP/IP stack. Use the client name with the **-z** option to direct any of the other **ipsec** commands to a specific client.

Using the pasearch command

You can use the **pasearch** commands listed in Table 71 to display information about the IPSec policy loaded by the Policy Agent for the stack:

Table 71. **pasearch** *commands*

Command	Description
pasearch -v a	Displays all IPSec policy
pasearch -v f	Displays IpFilterPolicy
pasearch -v k	Displays KeyExchangePolicy
pasearch -v l	Displays LocalDynVpnPolicy

The -p option can be used to obtain policy for a specific stack. Additional pasearch options can be used to obtain a more condensed display. See z/OS Communications Server: IP System Administrator's Commands for a complete description of the **pasearch** command syntax.

Using syslog messages

The IKE daemon uses syslogd to write informational messages to the local4 facility. These messages contain the following information:

- The state of the IKE daemon
- Successful and unsuccessful phase 1 and phase 2 negotiations
- Information about phase 1 and phase 2 negotiation failures

Additional IKE daemon debug information can be enabled by setting the IkeSyslog and PagentSyslogLevel parameters in the IKE configuration file. See “IKE daemon debug information” on page 350 for more details and sample IKE daemon syslog output.

If you are using the IBM Configuration Assistant for z/OS Communications Server to configure, the IKE Syslog level and the Policy Agent API Syslog level can be set from the IPSec: IKE Daemon Settings panel.

The Defense Manager daemon (DMD) uses syslogd to write informational messages to the local4 facility. DMD debug information can be enabled by setting the SyslogLevel parameter in the DMD configuration file. For more information, see “Defense Manager daemon debug information” on page 727.

The stack uses the TRMD daemon to write informational messages. The TRMD daemon uses syslogd to write these messages to the local4 facility. To enable many of these messages, IP filter logging must be turned on at both an IP filter policy level and an individual filter rule level. See “Steps for verifying IP security policy or defensive filter enforcement” on page 749 for details about enabling IP filter logging.

Chapter 31. Diagnosing OMPROUTE problems

This topic provides information and guidance to diagnose OMPROUTE problems, and contains the following subtopics:

- "Overview"
- "Definitions" on page 762
- "Diagnosing OMPROUTE problems" on page 762
- "OMPROUTE traces and debug information" on page 775
- "Starting OMPROUTE tracing and debugging from an MVS cataloged procedure or AUTOLOG" on page 776
- "TCP/IP services component trace for OMPROUTE" on page 786
- "Commands to enable, disable, and display the status of the OMPROUTE CTRACE" on page 790

Overview

For IPv4, OMPROUTE implements the Open Shortest Path First (OSPF) protocol described in RFC 1583, "OSPF Version 2" as well as the Routing Information Protocols (RIP) described in RFC 1058, "Routing Information Protocol" (RIP Version 1) and in RFC 1723, "RIP Version 2–Carrying Additional Information" (RIP Version 2).

For IPv6, OMPROUTE implements the IPv6 OSPF protocol described in RFC 2740, "OSPF for IPv6", as well as the IPv6 RIP protocol described in RFC 2080, "RIPng for IPv6".

OMPROUTE provides an alternative to the static TCP/IP BEGINROUTES definitions. When configured properly, the MVS host running with OMPROUTE becomes an active OSPF or RIP router in a TCP/IP network. The dynamic routing protocols are used to dynamically maintain the host routing table. For example, OMPROUTE can determine that a new route has been created, that a route is temporarily unavailable, or that a more efficient route exists.

OMPROUTE has the following characteristics:

- It is a z/OS UNIX application. It requires the z/OS UNIX file system to operate.
- OMPROUTE can be started from an MVS procedure, from the z/OS shell, or from AUTOLOG. See z/OS Communications Server: IP Configuration Guide for information about OMPROUTE.
- The OMPROUTE subagent provides an alternative to DISPLAY commands for displaying IPv4 Open Shortest Path First (OSPF) protocol configuration and state information. The subagent implements the Management Information Base (MIB) variables defined in Request for Comment (RFC) 1850. The OMPROUTE subagent is controlled by statements in the OMPROUTE configuration file. For details, see z/OS Communications Server: IP Configuration Reference.
- OMPROUTE needs to be started by a RACF authorized user ID.
- OMPROUTE needs to be in an APF authorized library.
- A one-to-one relationship exists between an instance of OMPROUTE and a TCP/IP stack. OSPF/RIP support on multiple TCP/IP stacks requires multiple instances of OMPROUTE.

© Copyright IBM Corp. 2000, 2015

- All IPv4 dynamic routes are deleted from the routing table upon initialization of OMPROUTE if there are IPv4 interfaces configured to OMPROUTE as RIP or OSPF interfaces.
- All IPv6 dynamic routes (with the exception of routes learned using the IPv6 Router Discovery protocol) are deleted from the routing table upon initialization of OMPROUTE if there are IPv6 interfaces configured to OMPROUTE as OSPF or RIP interfaces.
- IPv4 Internet Control Message Protocol (ICMP) redirects are ignored when OMPROUTE is active and there are IPv4 interfaces configured to OMPROUTE as RIP or OSPF interfaces.
- IPv6 ICMP redirects are ignored when OMPROUTE is active and there are IPv6 interfaces configured to OMPROUTE as OSPF or RIP interfaces.
- OMPROUTE does not make use of the BSD Routing Parameters. Instead, the maximum transmission unit (MTU), subnet mask, and destination address parameters for IPv4 interfaces are configured using the OSPF_Interface, RIP_Interface, and Interface statements in the OMPROUTE configuration file. Also, for IPv6, OMPROUTE does not update the stack's MTU sizes but learns them from the stack instead.

  Restriction: If using NCPROUTE, the BSD routing parameters in the BSDROUTINGPARMS TCP/IP configuration statement must be defined for the host-to-NCP channel interfaces, and the parameter values must match the corresponding values on the RIP_INTERFACE or INTERFACE statements in the OMPROUTE configuration file; otherwise, connection problems occur between NCPROUTE and its NCP clients.
- OMPROUTE uses the MVS operator console, SYSLOGD, STDOUT, and CTRACE for its logging and tracing.
  - The MVS operator console and SYSLOGD are used for major events such as initialization, termination, and error conditions.
  - STDOUT and z/OS UNIX file system files are used for detailed tracing and debugging.
  - CTRACE is used for the following purposes:
    - Tracing the receipt and transmission of OSPF/RIP packets
    - Tracing subagent/SNMP agent packets
    - Tracing communication between OMPROUTE and the TCP/IP stack
    - Detailed tracing and debugging

    For details on using TCP/IP Services Component trace support with OMPROUTE, see "TCP/IP services component trace for OMPROUTE" on page 786 and Chapter 5, "TCP/IP services traces and IPCS support," on page 45.
- If you want to communicate a routing protocol over an interface, configure the interface to OMPROUTE using the OSPF_INTERFACE, RIP_INTERFACE, IPV6_OSPF_INTERFACE, or IPV6_RIP_INTERFACE configuration statement.
- IPv4 interfaces that are not involved in the communication of the RIP or OSPF protocol (except VIPA interfaces) must be configured to OMPROUTE using the INTERFACE configuration statement, unless it is a non-point-to-point interface and all default values are acceptable as specified on the INTERFACE statement. All IPv4 interfaces known to the TCP/IP stack should be defined to OMPROUTE with the correct subnet mask and MTU values. For IPv4 interfaces that are not defined to OMPROUTE, OMPROUTE assigns default subnet mask and MTU values to the interfaces, with possibly undesirable results.

- IPv6 interfaces that are not involved in the communication of the OSPF or RIP protocol defaults to IPv6 generic interfaces when Global_Options Ignore_Undefined_Interfaces is coded to No (default value). The IPv6_Interface statement can be used if the IPv6 (generic) interface default values are not acceptable or you want to define additional IPv6 prefixes on the IPv6_Interface statement. If Global_Options Ignore_Undefined_Interfaces is coded to Yes, code IPv6_INTERFACE statements for all IPv6 Interfaces not involved in communication of OSPF or RIP that you want OMPROUTE to recognize.
- OMPROUTE uses a standard message catalog. The message catalog must be in the z/OS UNIX file system. The directory location for the message catalog path is set by the environment variables NLSPATH and LANG.
- If you want OMPROUTE to completely ignore IPv4 and IPv6 interfaces that are not defined to it, code the GLOBAL_OPTIONS statement with IGNORE_UNDEFINED_INTERFACES=YES in the OMPROUTE configuration file. For details, see z/OS Communications Server: IP Configuration Guide.
- OMPROUTE is enhanced with Virtual IP Addressing (VIPA) to handle network interface failures by switching to alternate paths. The virtual routes are included in the OSPF and RIP advertisements to adjacent routers. Adjacent routers learn about virtual routes from the advertisements and can use them to reach the destinations at the MVS host.
- OMPROUTE allows for the generation of multiple, equal-cost routes to a destination, thus providing load-balancing support.
- During a temporary shortage in storage, such as CSM ECSA or CSM data space conditions or reaching the TCP/IP defined limits for ECSA or private storage, OMPROUTE temporarily suspends the route timeout processing. This is done in an attempt to prevent the loss of routing information about the local host.

OMPROUTE works best without non-replaceable static routes, and the use of non-replaceable static routes (defined using the BEGINROUTES TCP/IP configuration statement) is not recommended. Non-replaceable static routes might interfere with the discovery of a better route to the destination as well as inhibit the ability to switch to another route if the destination should become unreachable by way of the static route. For example, if you define a non-replaceable static host route through one interface and that interface becomes unreachable, OMPROUTE does not define a route to that same host through an alternate interface.

If you must define static routes, all static routes are considered to be of equal cost and non-replaceable static routes are not replaced by OSPF or RIP routes. Use extreme care when working with static routes and OMPROUTE. Set IMPORT_STATIC_ROUTES = YES on the AS_Boundary Routing or IPv6_AS_Boundary_Routing configuration statement, or both. Alternatively, set SEND_STATIC_ROUTES = YES on the RIP_Interface or IPv6_RIP_Interface configuration statement, or both. This allows the static routes to be advertised to other routers.

You can define static routes as replaceable. Unlike non-replaceable static routes, replaceable static routes are always replaced by dynamic routes learned by OMPROUTE. In other words, a replaceable static route is used only if no dynamic route is known to the destination. Replaceable static routes can be thought of as last resort routes to reach a destination when no dynamic route is known.

Definitions

OMPROUTE must be defined correctly to TCP/IP. For detailed information about TCP/IP definitions, see the information about configuring OMPROUTE in z/OS Communications Server: IP Configuration Reference.

Diagnosing OMPROUTE problems

Problems with OMPROUTE are generally reported under one of the following categories:

- Abends
- OMPROUTE connection problems
- Routing failures
- Adjacency failures
- Malformed and unexpected packets

These categories are described in the following topics.

Abends

An abend during OMPROUTE processing should result in messages and error-related information being sent to the system console. A dump of the error is needed unless the symptoms match a known problem. If a dump was not taken, ensure the Language Environment run-time options TRAP(ON,NOSPIE) TERMTHDACT(UAIMM) are set for OMPROUTE.

OMPROUTE connection problems

OMPROUTE connection problems are reported when OMPROUTE is unable to connect to TCP/IP or to one of the ports required for OSPF or RIP communication. These problems are generally caused by an error in the configuration or definitions in TCP/IP.

In a common INET environment (multiple stacks), OMPROUTE attempts to connect to a stack whose name is determined by the TCPIPjobname keyword in the resolver configuration data set or file. If OMPROUTE cannot determine the TCPIPjobname, it uses a default of INET. If OMPROUTE cannot communicate with the stack pointed to by TCPIPjobname or is unable to initialize its required ports, it issues an error message describing the problem and then terminates.

For details about diagnosing problems while attempting to connect to the SNMP agent, see “SNMP connection problems” on page 614.

Routing failures

Routing problems are usually the result of outages in a network and a lack of alternative routing paths available for recovery. See “Steps for verifying IP routing to a destination when not using policy-based routing” on page 29 and “Steps for diagnosing problems with IP routing to a destination when using policy-based routing” on page 31 for help with diagnosing routing failures

Table 72 on page 763 describes command terms used in this section.

Table 72. OMPROUTE command terms

Term	Description
NETSTAT ROUTE	Refers to the Netstat ROUTE/-r command and the netstat route commands used on other platforms.
OMPROUTE RTTABLE	Refers to the D TCPIP,tcpipjobname,OMPROUTE,RTTABLE command for displaying OMPROUTE IPv4 route tables.
OMPROUTE RT6TABLE	Refers to the D TCPIP,tcpipjobname,OMPROUTE,RT6TABLE command for displaying OMPROUTE IPv6 route tables.
PING	Refers to z/OS UNIX **ping**, TSO PING, and the **ping** commands used on other platforms.
Traceroute	Refers to z/OS UNIX **traceroute**, TSO TRACERTE, and the **traceroute** commands used on other platforms.

Analyzing routing failures

Guidelines: When analyzing routing failures, follow these guidelines:

- Make sure that the address used in attempting to contact the remote host is a valid IP address.
- Make sure routing is possible in both directions. For most TCP/IP communication, two-way routing is required. The origin must have routes to reach the destination, and the destination must have routes to reach the origin. If NETSTAT ROUTE at the origin shows correct routing, you must also use NETSTAT ROUTE at the destination to verify that it can send replies back to the origin. If there are intermediate hops between the source and destination, all routing tables must have routing information. For example, if the origin node routing table indicates that the first hop to reach the destination is router A, then the router A routing table must also have a valid, active route to the destination, and so on. This also applies to the return route.
- Also, this is affected by SOURCEVIPA. If SOURCEVIPA is enabled at the origin of the communication, then the destination and all intermediate hops must be able to route back to the VIPA.
- If the NETSTAT ROUTE output on the source, the destination, or an intermediate hop does not show the expected routes, do one or more of the following:
  - Make sure that the routers involved in providing routing information are operational and participating in the correct routing protocol.
  - Make sure that the necessary physical connections are active.
  - Use the OMPROUTE DISPLAY commands described in z/OS Communications Server: IP System Administrator's Commands to determine whether anything in the configuration or current state of OMPROUTE has caused the unexpected NETSTAT ROUTE information.

Documenting routing failures

You should gather documentation described in "Documentation for the IBM Support Center" on page 41 for initial diagnosis of all routing failures. If dynamic

routing is being provided by OMPROUTE and the expected dynamic routes have not been installed in the stack route table, the following documentation should also be available:

- MVS system log
- SYSLOGD
- The data set containing OMPROUTE trace and debug information. If OMPROUTE trace and debug information is being redirected to the OMPROUTE CTRACE internal buffer, this buffer is included in a dump of the OMPROUTE address space. For details, see "OMPROUTE traces and debug information" on page 775 and "TCP/IP services component trace for OMPROUTE" on page 786.
- Output from OMPROUTE RTTABLE or RT6TABLE commands. If using policy-based routing, collect output for the appropriate route tables.
- Output from any other OMPROUTE DISPLAY commands used.

Adjacency failures

OMPROUTE adjacency failures are reported when OMPROUTE is unable to establish adjacency with a neighboring router, or loses an established adjacency, over one of its network interfaces. The following error messages are used to report adjacency failures:

```
For IPv4: EZZ7921I OSPF Adjacency Failure, neighbor neighbor, old state state,
 new state state, event event

For IPv6: EZZ7954I IPv6 OSPF Adjacency Failure, neighbor neighbor, old state ostate,
new state nstate, event event
```

In addition to the adjacency failure message, if the futile neighbor state loop detection is enabled in OMPROUTE (Max_Adj_Attempt parameter on OSPF and IPV6OSPF statements), the following error messages will be issued to report futile neighbor state loops for the adjacency attempts:

```
For IPv4: EZZ8157I jobname IPv4 OSPF detected futile neighbor state loop with neighbor neighbor
on interface interface after threshold_value adjacency attempts

For IPv6: EZZ8157I jobname IPv6 OSPF detected futile neighbor state loop with neighbor neighbor
on interface interface after threshold_value adjacency attempts
```

If OMPROUTE is configured with redundant parallel interfaces (primary and backup) attached to the same LAN segment, OMPROUTE will try to form adjacency with a neighboring designated router over the alternate redundant interface when the futile neighbor state loop has been detected on the problematic interface. The following informational messages will be issued to report the interface changes::

```
For IPv4: EZZ8158I jobname IPv4 OSPF could not establish adjacency on interface interface1 -
attempting to establish adjacency on interface interface2

For IPv6: EZZ8158I jobname IPv6 OSPF could not establish adjacency on interface interface1 -
attempting to establish adjacency on interface interface2
```

For information about the Max_Adj_Attempt parameter, see the OMPROUTE topic in z/OS Communications Server: IP Configuration Reference.

For information about futile neighbor state loops, see the topic about network design considerations with z/OS Communications Server in z/OS Communications Server: IP Configuration Guide.

OMPROUTE adjacency failures can also occur during the processing of received OSPF packets that are malformed or unexpected. For example, a neighboring router can be sending an incorrectly built Link State Advertisement (LSA) or sending an unexpected AS External LSA for a stub area to OMPROUTE. For more information, see "Malformed and unexpected packets" on page 773.

Analyzing adjacency failures

An adjacency failure is reported by the neighbor event code in the OMPROUTE informational message provided. The events that are associated with adjacency failures are indicated by codes 7 – 15:

Table 73. Neighbor event codes

Event	Description	Explanation
1	Hello received	OMPROUTE received hello packets from a neighbor.
2	Start sending hellos	OMPROUTE is ready to start sending hello packets to a neighbor.
3	Two-way communication	OMPROUTE reached bidirectional communication with a neighbor. The neighbor and OMPROUTE received and acknowledged hello packets from each other.
4	Ready to form adjacency	OMPROUTE is ready to establish adjacency with a neighbor.
5	Master/slave role negotiation is done	OMPROUTE completed negotiating the master and slave roles with a neighbor. This signals the start of sending and receiving of database descriptor packets.
6	Database exchange is done	OMPROUTE completed exchanging the database descriptor packets for the network topology information. Each router now knows which part of its link state database (LSDB) is outdated.
7	Sequence number mismatch	OMPROUTE received a sequence number mismatch in a database descriptor packet. A neighbor might be attempting to restart the adjacency for some reason, resulting in sequence number mismatches. This event indicates that the neighbor was not receiving hello packets from OMPROUTE, experienced event 12 on its side, and is trying to restart.
8	Bad link state request	OMPROUTE received a bad Link State request (LSA) packet from a neighbor. A subverted router in the network possibly modified the contents of an LSA to result in maximum sequence number or maximum age attacks. From the LSA floods, the neighbors replace the good LSA with the bad LSA as newer in their databases until they naturally age out in a maximum of 1 hour.
9	Loading is done	OMPROUTE completed the loading of the link state information for the link state database, which is based on the link state request packets from a neighbor.

Table 73. Neighbor event codes (continued)

Event	Description	Explanation
10	One-way communication only	OMPROUTE does not see its router ID in the hello packets that are received by a neighbor. If a neighbor is not receiving the hello packets with the OMPROUTE router ID, it assumes that OMPROUTE is down and removes the OMPROUTE router ID from its list of neighbors. The neighbor does not include the OMPROUTE router ID in the hello packets sent to OMPROUTE.
11	Neighbor is down	Because the OMPROUTE routing interface is down, it is unable to communicate with a neighbor over that interface. The neighbor is set to down state.
12	No hellos were seen recently	OMPROUTE did not receive hello packets from a neighbor for a full DEAD_ROUTER_INTERVAL and as a consequence, OMPROUTE assumes that the neighbor is down.
14	Start adjacency establishment	OMPROUTE starts to establish an adjacency with a neighbor after it reaches bidirectional communication with the neighbor. OMPROUTE sends database descriptor packets to negotiate master or slave status with the neighbor to establish an adjacency.
15	Failure to thrive	OMPROUTE was trying to establish adjacency with a neighbor but failed to complete within the DB_EXCHANGE_INTERVAL
16	Adjacency attempts threshold is reached	OMPROUTE has reached the futile neighbor state loop threshold (DR_MAX_ADJ_ATTEMPT) for adjacency attempts with a neighbor. For more information about the actions to be taken, see Preventing futile neighbor state loops during adjacency formation in z/OS Communications Server: IP Configuration Guide.

With the neighbor event code, the old neighbor state indicates the highest neighbor state that OMPROUTE reached for the adjacency attempt and the new neighbor state indicates the changed state. When OMPROUTE cannot reach the two-way state (8) or loses the bidirectional communication with a neighboring router, OMPROUTE reverts to the lesser neighbor states 1 – 4. On non-multiaccess networks such as point-to-point links, OMPROUTE attempts to reach full adjacency with all neighboring routers. On multiaccess networks such as LANs (including HiperSockets LANs), OMPROUTE attempts to reach full adjacency only with neighbors that are designated routers or backup designated routers. If OMPROUTE is the designated router or the backup designated router on a multiaccess network, it attempts to reach full adjacency with all neighbors.

Guideline: OMPROUTE is not intended to be a designated router or a backup designated router on multiaccess networks that include dedicated routers. The dedicated routers can provide the designated router function. Configure

OMPROUTE to assume the designated router role only when no dedicated routers are available, such as in a HiperSockets LAN in which all the neighbors are z/OS LPARs.

Starting at the two-way state for bidirectional communication, the neighbor states 8 – 128 are used by OMPROUTE in attempts to reach full adjacency with a neighboring router, as shown in the following table:

Table 74. Neighbor states

State	Description	Explanation
1	Neighbor is down or unreachable	No hello packets are received by the neighbor. Possible causes are as follows: • The routing interface in OMPROUTE is down • The neighbor is inactive • The neighbor is not OSPF-enabled • The neighbor is not sending hello packets with its unique router ID and matching attributes (hello interval, dead router interval, subnet mask, area ID, and security scheme) • Hello packets were dropped between OMPROUTE and the neighbor • Hello packets are corrupted
2	Attempting to contact neighbor	OMPROUTE sent hello packets to the network but no hello packets were received from a neighbor. This state is valid only for manually configured neighbors in a non-broadcast multiaccess (NBMA) environment. For possible causes, see state 1.
4	One-way communication	OMPROUTE received hello packets from a neighbor with the router ID of the neighbor. If OMPROUTE remained in this state, acknowledgments to hello packets that contain the OMPROUTE router ID were not received from the neighbor. Possible causes are as follows: • The neighbor did not receive the hello packets from OMPROUTE that contain its router ID • The neighbor is not including the OMPROUTE router ID in its hello packets that are sent to OMPROUTE • Hello packets were dropped between OMPROUTE and the neighbor • Hello packets are corrupted • New designated router must be reelected after lost communications with designated router and backup designated router

Table 74. Neighbor states (continued)

State	Description	Explanation
8	Two-way communication	The neighbor and OMPROUTE received and acknowledged hello packets from each other. In the list of neighbors, OMPROUTE has the neighbor router ID and the neighbor has the OMPROUTE router ID as learned from the hello packets. If OMPROUTE remained in or reverted to this state during the adjacency formation, possible causes are as follows: • Neither OMPROUTE or the neighbor is the designated router or backup designated router for the network. On a multiaccess network, full adjacencies are established only with the designated router and its backup. • A database descriptor packet is received with: – Mismatched sequence number – Unexpected `init` bit set – `Options` field differs from the last `Options` field received • Acknowledgments to database descriptor packets were dropped between OMPROUTE and the neighbor.
16	Database exchange start	OMPROUTE is negotiating master and slave roles with a neighbor. Designated routers and backup designated routers establish a master-slave relationship and choose the initial sequence number for the adjacency formation. The neighbor with the higher router ID becomes the master and starts the exchange. The master is also responsible for the sequence number increment. If OMPROUTE remained in this state, possible causes are as follows: • Mismatched interface maximum transmission unit (MTU) values between OMPROUTE and the neighbor result in packet losses • Duplicate router IDs on neighbors • Access list is blocking the unicasts • NAT translating the unicasts

Table 74. Neighbor states (continued)

State	Description	Explanation
32	Database exchange	OMPROUTE is exchanging database information with a neighbor in the form of database descriptor packets. These packets contain LSA headers only and describe the contents of the entire link state database (LSDB). Each database descriptor packet has a sequence number that only the master can increment and that the slave explicitly acknowledges. The packet contents that are received are compared with the information contained in the LSDB for new or more current link state information. Routers also send link state request packets and link state update packets that contain the entire LSA. If OMPROUTE remained in this state, possible causes are as follows: • Corrupted database descriptor packets that are sent by a neighbor or by a network switch • See state 16 for other possible causes
64	Loading	OMPROUTE is requesting newer pieces from a neighbor database that are more up-to-date in the form of link state request packets. Based on the information that the database descriptor packets provide, routers send link state request packets. The neighbor then provides the requested link state information in link state update packets. During the adjacency formation, if a router receives an outdated or missing LSA, it requests that LSA by sending a link state request packet. All link state update packets are acknowledged. If OMPROUTE remained in this state, possible causes are as follows: • Corrupted link state request packets that are sent by a neighbor or by a network switch • See state 16 for other possible causes
128	Full	OMPROUTE established full adjacency with a neighbor. Routers achieve the full state with their designated router and backup designated router only, and neighbors always see each other as two-way state (8).

The order of the neighbor transit states for establishing adjacency is 8, 16, 32, 64, and 128. Whenever a problem is detected for some reason between those states before full adjacency is reached, OMPROUTE resets the neighbor state to two ways (8) and repeats the process on a continuous basis even to the point where it becomes futile. A futile neighbor state loop is seen as a successive repetitive pattern of transit states and ones that do not seem to reach full adjacency. For example, typical patterns are: 8-16, 8-16 or 8-16-32, and 8-16-32. After each adjacency failure, OMPROUTE continues to attempt to establish adjacency with a

neighbor over the same network interface. If futile neighbor state loop detection is enabled and if there are redundant parallel interfaces (primary or backup) attached to the same LAN segment available, OMPROUTE suspends the problematic interface and tries the adjacency attempt over the alternative interface again. The other option is to use the MODIFY OMPROUTE commands to manually suspend and activate an alternative redundant parallel OSPF interface so that adjacency with the neighbor is attempted over that interface.

For information about futile neighbor state loops, see the topic about network design considerations with z/OS Communications Server in z/OS Communications Server: IP Configuration Guide. For details about MODIFY OMPROUTE commands, see z/OS Communications Server: IP System Administrator's Commands.

OMPROUTE can drop adjacencies under the following conditions:

- Other workloads on the z/OS system keep OMPROUTE from dispatching enough processor cycles:
  - Dumps are being taken while OMPROUTE is running. All address spaces are marked non-dispatchable during a dump processing. If the dump takes longer than a DEAD_ROUTER_INTERVAL, the adjacencies fail.
  - There are too many other address spaces that are running at a higher priority than OMPROUTE. Because OMPROUTE is a time-sensitive application and manages the TCP/IP routing table, set the OMPROUTE priority to be one less than the TCP/IP dispatching priority. If you are using WLM goal modes, place OMPROUTE in the same service class as TCP/IP.
- Not enough dispatching for OMPROUTE:
  - The dispatching priority for OMPROUTE is too low. Either increase the OMPROUTE dispatching priority or increase the DEAD_ROUTER_INTERVAL values.
  - OMPROUTE is running as a BPXBATCH program. Because there are other applications that are using the BPXBATCH program, they might steal processor cycles from OMPROUTE. Change OMPROUTE not to use the BPXBATCH program.
- Increased workload in OMPROUTE:
  - OMPROUTE is a designated router or a backup designated router, and link state database management-related tasks can contribute to high workloads. These high workloads can affect OMPROUTE in processing of inbound hello packets necessary to maintain adjacencies with its neighbors. Either increase the DEAD_ROUTER_INTERVAL values or change the ROUTER_PRIORITY values to reduce the likelihood of OMPROUTE becoming elected as a designated router. A z/OS system is not designed to be a full-fledged router and it is best to offload this work of link state database management to the neighboring network routers. That is, configure the network routers on the attached LAN segment to be elected as designated routers when possible.
  - OMPROUTE is running with too much tracing. OMPROUTE debug trace with file I/O can contribute to adjacency failures (for example, missed hello packets). Use the OMPROUTE CTRACE method when possible.
  - OMPROUTE routing table is large. After OMPROUTE has more than 1000-2000 routes, adjacency failures (for example, missed hello packets) might occur because of the increased workload from processing routing table and link state updates. Configure OMPROUTE to use stub areas when possible and try to keep z/OS out of backbone areas.

  - OMPROUTE has too many adjacencies. Too many adjacencies might be notable when you are using XCF in a sysplex environment. Because of increased workload from adjacency communications, adjacency failures (for example, missed hello packets) might occur. Determine whether it is necessary for XCF interfaces to be configured to use OSPF.
- Network hardware problem:
  - Attached switch or router not functioning correctly.
  - Poor or faulty network cable connections.

    A network hardware problem that is beyond detection by TCP/IP or OMPROUTE can contribute to adjacency failures and futile neighbor state loops. If futile neighbor state loop detection is enabled and if there are redundant parallel interfaces that are attached to the same LAN segment, OMPROUTE attempts adjacency with the neighbor over an alternative interface. When necessary, use the MODIFY OMPROUTE commands to manually suspend and activate an alternative redundant parallel OSPF interface so that adjacency with the neighbor is attempted over that interface. OMPROUTE might circumvent the network hardware problem by using the alternative interface.

For the symptom of missed inbound or outbound hello packets, the TCP/IP stack might not be getting dispatched often enough to forward the hello packets to or from OMPROUTE. In this case, ensure that appropriate dispatching priorities are assigned to the TCP/IP stack and OMPROUTE.

To track adjacency problems, take the following steps:

- Issue the command to display the OSPF interfaces and analyze the following fields from the report:
  - STATE for the current interface state
  - #NBRS for the total number of neighbors whose hellos were received, plus those that were configured.
  - #ADJS for the total number of neighbors in state exchange or greater. These neighbors are the neighbors with whom the router is synchronized or is in the process of synchronization.
- Issue the command to display a detailed OSPF interface and analyze the following fields from the report:
  - DESIGNATED ROUTER to determine whether OMPROUTE is a designated router or not.
  - BACKUP DR to determine whether OMPROUTE is a backup designated router or not.
  - DR PRIORITY to determine the interface router priority. A higher value indicates that this OMPROUTE is more likely to become the designated router. A value of 0 indicates that OMPROUTE can never become the designated router.
  - #NEIGHBORS for total number of neighbors whose hellos were received, plus those that were configured.
  - #ADJACENCIES for total number of neighbors in state Exchange or greater. These neighbors are the neighbors with whom the router is synchronized or is in the process of synchronization.
  - #FULL ADJS for total number of full adjacencies. This number is the number of neighbors whose state is Full (and therefore with which the router synchronized databases).

  - #MCAST FLOODS for the total number of link state updates that flooded the interface (not counting retransmissions).
- Issue the command to display the OSPF neighbors and analyze the following field from the report:
  - STATE for the current neighbor state.
- Issue the command to display a detailed OSPF neighbor and analyze the following fields from the report:
  - NEIGHBOR STATE for the current neighbor state.
  - DR PRIORITY to determine the neighbor router priority.
  - #ADJ RESETS for the total number of transitions to state ExStart from a higher state.
  - #NBR LOSSES for total number of times the neighbor made the transition to the down state.

Documenting adjacency failures

The following documentation should be available for initial diagnosis of adjacency failures:

- MVS system log
- SYSLOGD
- The data set containing OMPROUTE trace and debug information, unless trace and debug information is being redirected to the OMPROUTE CTRACE internal buffer, which is automatically included in a dump of the OMPROUTE address space. For details, see "OMPROUTE traces and debug information" on page 775.
- TCP/IP and OMPROUTE CTRACE. For information about generating an OMPROUTE Component Trace, see "TCP/IP services component trace for OMPROUTE" on page 786.
- Output from appropriate OMPROUTE DISPLAY commands as described in z/OS Communications Server: IP System Administrator's Commands.

  Use the DISPLAY TCPIP,,OMROUTE,OPTIONS command to display the global configuration options.

  If using IPv4, use the following commands to display OSPF configuration, interfaces, neighbors, and routing table:
  - DISPLAY TCPIP,,OMPROUTE,OSPF,LIST,ALL
  - DISPLAY TCPIP,,OMPROUTE,OSPF, INTERFACES
  - DISPLAY TCPIP,,OMPROUTE,OSPF, INTERFACE,NAME=if_name
  - DISPLAY TCPIP,,OMPROUTE,OSPF,NEIGHBOR
  - DISPLAY TCPIP,,OMPROUTE,OSPF,NEIGHBOR,IPADDR=ip-addr
  - DISPLAY TCPIP,,OMPROUTE,RTTABLE

  If using IPv6, use the following commands to display OSPF configuration, interfaces, and neighbors, and routing table:
  - DISPLAY TCPIP,,OMPROUTE,IPV6OSPF,ALL
  - DISPLAY TCPIP,,OMPROUTE,IPV6OSPF, INTERFACES
  - DISPLAY TCPIP,,OMPROUTE,IPV6OSPF, INTERFACE,NAME=if_name [Note: ID=if-id can be used instead of NAME parameter]
  - DISPLAY TCPIP,,OMPROUTE,IPV6OSPF,NEIGHBOR
  - DISPLAY TCPIP,,OMPROUTE,IPV6,NEIGHBOR,ID=router-id [Note: IFNAME=if_name can be used instead of ID parameter]
  - DISPLAY TCPIP,,OMPROUTE,RT6TABLE

- When applicable, the dumps of TCPIP and VTAM address spaces. A SLIP on OMPROUTE adjacency failing message (for example, EZZ7921I) may be used to capture the OMPROUTE trace and debug information as well as the dumps of TCPIP and VTAM address spaces at the time of the error.

Malformed and unexpected packets

An OSPF packet is considered malformed and is discarded if any of the following are present in the packet:

- Incorrect version in the OSPF packet header.
- Incorrect packet type in the OSPF header.
- Incorrect length value in the OSPF packet header.
- Incorrect checksum in the OSPF packet header.
- Incorrect LSA type in an LSA header.
- Incorrect length value in an LSA header.
- Incorrect checksum in an LSA header.
- Incorrect LSA age field in an LSA header.
- Incorrect LSA scope value in an LSA header.
- Incorrect prefix length value in an LSA that includes address prefixes.
- Incorrect value for the number of entries in a packet or LSA. For example, a bad number of LSAs value in a Link State Update packet, a bad number of links value in a Router LSA, or a bad number of prefixes value in a Link LSA or Intra-Area Prefix LSA.

An OSPF packet is considered unexpected and is discarded if any of the following are present in the packet:

- Unexpected AS External LSA received over an interface that is attached to an OSPF stub area.
- Unexpected AS External LSA received over a virtual link.
- Unexpected AS Scope LSA received over an interface that is attached to an OSPF stub area.

Analyzing malformed and unexpected packets

OMPROUTE issues console messages to report errors in the received OSPF packets. In addition, OMPROUTE tracks the number of malformed and unexpected packets that are received for the OSPF packet statistics. If such packets are being received, contact the manufacturer of the routing application that is sending them. The malformed packets are viewable for analysis from the following documentation:

- Network sniffer trace
- IP packet trace
- OMPROUTE trace and debug output

An IP packet trace (SYSTCPDA) can be used to obtain traces of IPv4 and IPv6 OSPF packets flowing from and into TCP/IP on a z/OS Communications Server host. The PKTTRACE statement or command enable you to trace IP packets as they enter or leave TCP/IP. The IP packets can be filtered to the OSPF protocol number by using the PROT option on the PKTTRACE statement or command.

To track the number of malformed and unexpected packets that are received, take the following steps:

- Issue the command to display detailed information about an OSPF interface and analyze the # ERR PKTS RCVD field in the report to see the total number of packets received on the interface that contained errors.
- Issue the command to display detailed information about an OSPF neighbor and analyze the # ERR LS RVCD field in the report to see the total number of advertisements received from the neighbor that were unexpected or contained errors.

Documenting malformed and unexpected packets

If the manufacturer of the routing application cannot provide documentation that indicates that it is sending the malformed or unexpected packets that are being received by OMPROUTE , the following documentation is available for initial diagnosis.

- A network sniffer trace.
- The data set that contains the IP packet trace (SYSTCPDA), unless the packet trace is being redirected to the TCP/IP address space. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for more information.
- The data set that contains OMPROUTE trace and debug information. If OMPROUTE trace and debug information is being redirected to the OMPROUTE CTRACE internal buffer, this buffer is included in a dump of the OMPROUTE address space. See "OMPROUTE traces and debug information" on page 775 for more information.
- Output from the appropriate OMPROUTE DISPLAY commands as described in z/OS Communications Server: IP System Administrator's Commands.

If the malformed or unexpected packets that are being received are IPv4 OSPF packets, use the following commands to display your IPV4 OSPF interfaces and neighbors:

- DISPLAY TCPIP,,OMPROUTE,OSPF, INTERFACES
- DISPLAY TCPIP,,OMPROUTE,OSPF, INTERFACE,NAME=*if_name*
- DISPLAY TCPIP,,OMPROUTE,OSPF,NEIGHBOR
- DISPLAY TCPIP,,OMPROUTE,OSPF,NEIGHBOR,IPADDR=*ip-addr*

If the malformed or unexpected packets that are being received are IPv6 OSPF packets, use the following commands to display your IPv6 OSPF interfaces and neighbors:

- DISPLAY TCPIP,,OMPROUTE,IPV6OSPF, INTERFACES
- DISPLAY TCPIP,,OMPROUTE,IPV6OSPF, INTERFACE,NAME=*if_name*

  Note: ID=*if-id* can be used instead of NAME=*if_name*
- DISPLAY TCPIP,,OMPROUTE,IPV6OSPF,NEIGHBOR
- DISPLAY TCPIP,,OMPROUTE,IPV6OSPF,NEIGHBOR,ID=*router-id*

  Note: Specify IFNAME=*if_name* if the neighbor that is specified by the ID parameter has more than one neighbor relationship with OMPROUTE (for example if there are multiple IPv6 OSPF interfaces that connect OMPROUTE to the neighbor).

When applicable, you can use a SLIP on the OMPROUTE error message that is reporting the receipt of the malformed or unexpected packet to capture the OMPROUTE trace and debug information and the dump of the TCP/IP address space that contains the IP packet trace at the time of the error.

OMPROUTE traces and debug information

There are many TCP/IP traces that can be useful in identifying the cause of OMPROUTE problems. OMPROUTE's use of the MVS Component Trace support is also useful (see “TCP/IP services component trace for OMPROUTE” on page 786). This section describes the OMPROUTE internal traces. OMPROUTE internal tracing and debugging can be started when OMPROUTE is started. Also, the MODIFY command can be used to start, stop, or alter OMPROUTE tracing and debugging after OMPROUTE has been started.

This topic describes each of these methods.

Starting OMPROUTE tracing and debugging from the z/OS UNIX System Services shell

If OMPROUTE is started from the z/OS UNIX System Services shell command line (using the **omproute** command), you can specify the following parameters to indicate the level of tracing or debugging that you want.

- -tn and -6tn (where n is a supported trace level)

  These options specify the OMPROUTE external tracing levels, with -tn covering both OMPROUTE initialization and IPv4 routing protocols and -6tn covering IPv6 routing protocols. These options provide information about the operation of the routing application and can be used for many purposes, such as debugging a configuration, education on the operation of the routing application, verification of test cases, and so on. The following trace levels are supported:

  – 1 = Informational messages

  – 2 = Formatted packet trace

- -sn (where n is a supported debug level)

  This option specifies the internal debugging level for the OMPROUTE subagent. It provides internal debugging information needed for debugging problems. The following level is supported:

  – 1 = Internal debugging messages. This turns on DPIdebug(2).

- -dn and -6dn (where n is a supported debug level)

  These options specify the OMPROUTE internal debugging levels, with -dn covering both OMPROUTE initialization and IPv4 routing protocols and -6dn covering IPv6 routing protocols. These options provide internal debugging information needed for debugging problems. The following levels are supported:

  – 1 = Internal debugging messages.

  – 2 = Unformatted hexadecimal packet trace

  – 3 = Function entry or exit trace

  – 4 = Task add or run

Guidelines:

- The -tn, -6tn, -dn, and -6dn options affect OMPROUTE performance. As a result, you might have to increase the Dead Router Interval on OSPF and IPv6 OSPF interfaces to prevent neighbor adjacencies from collapsing.
- The trace and debug levels are cumulative; each level includes all lower levels. For example, -t2 provides formatted packet trace and informational messages. You can enter more than one parameter by inserting a space after each parameter, for example, **omproute -t1 -d2**, which is the trace level most often requested by support. For more information, see APAR II12026.
- Parameters can be specified in mixed case.

Starting OMPROUTE tracing and debugging from an MVS cataloged procedure or AUTOLOG

The OMPROUTE tracing and debugging are controlled by parameters on PARM= when OMPROUTE is started from an MVS cataloged procedure or AUTOLOG. For example:

```
//OMPROUTE EXEC PGM=OMPROUTE,REGION=10M,TIME=NOLIMIT,
// PARM=('POSIX(ON) ENVAR("_CEE_ENVFILE=DD:STDENV")/-t2 -d1')
```

For a description of the parameters that can be specified, see "Starting OMPROUTE tracing and debugging from the z/OS UNIX System Services shell" on page 775.

Starting OMPROUTE tracing and debugging using the MODIFY command

Whether you start OMPROUTE from the z/OS UNIX System Services shell or from an MVS cataloged procedure, you can use the MODIFY command to start logging or tracing, to stop logging or tracing, and to change the level of logging or tracing.

The syntax for these MODIFY commands follows:

- MODIFY *procname*,TRACE=*trace-level*

  Use the TRACE command to change the trace level for OMPROUTE initialization as well as IPv4 routing protocols.
  - TRACE=0 turns off OMPROUTE tracing.
  - TRACE=1 gives all the informational messages.
  - TRACE=2 gives the informational messages plus formatted packet tracing.
- MODIFY *procname*,TRACE6=*trace-level*

  Use the TRACE6 command to change the trace level for IPv6 routing protocols.
  - TRACE6=0 turns off OMPROUTE tracing.
  - TRACE6=1 gives all the informational messages.
  - TRACE6=2 gives the informational messages plus formatted packet tracing.
- MODIFY *procname*,DEBUG=*debug-level*

  Use the DEBUG command to change the debug level for OMPROUTE initialization as well as IPv4 routing protocols.
  - DEBUG=0 turns off OMPROUTE debugging.
  - DEBUG=1 gives internal debug messages.
  - DEBUG=2 gives the same as DEBUG=1 plus hexadecimal packet tracing.
  - DEBUG=3 gives the same as DEBUG=2 plus module entry and exit.
  - DEBUG=4 gives the same as DEBUG=3 plus task add and run.
- MODIFY *procname*,DEBUG6=*debug-level*

  Use the DEBUG6 command to change the debug level for IPv6 routing protocols.
  - DEBUG6=0 turns off OMPROUTE debugging.
  - DEBUG6=1 gives internal debug messages.
  - DEBUG6=2 gives the same as DEBUG6=1 plus hexadecimal packet tracing.
  - DEBUG6=3 gives the same as DEBUG6=2 plus module entry and exit.
  - DEBUG6=4 gives the same as DEBUG6=3 plus task add and run.
- MODIFY *procname*,SADEBUG=*trace-level*

Use the SADEBUG command to start and stop message logging for the OMPROUTE subagent and to stop DPI tracing:

- SADEBUG=0 stops message logging for the OMPROUTE subagent and issues DPIdebug(0) to stop DPI tracing.
- SADEBUG=1 generates all messages by the OMPROUTE subagent and DPIdebug(2).

Destination of OMPROUTE trace and debug output

If the OMPROUTE CTRACE with option DEBUGTRC (or option ALL) is not enabled, then output from OMPROUTE tracing and debugging is written to the debug output destination. The debug output destination is based on the OMPROUTE_DEBUG_FILE and OMPROUTE_IPV6_DEBUG_FILE environment variables. If OMPROUTE was started without tracing enabled and OMPROUTE_DEBUG_FILE/OMPROUTE_IPV6_DEBUG_FILE is not defined and tracing is started later using the MODIFY command, the trace output destination is $TMP/omproute_debug, where $TMP is the value of the TMP environment variable.

When OMPROUTE_DEBUG_FILE is defined, the first trace file created for OMPROUTE initialization and IPv4 routing protocol tracing is named using the value coded on OMPROUTE_DEBUG_FILE. When OMPROUTE_IPV6_DEBUG_FILE is defined, the first trace file created for IPv6 routing protocol tracing is named using the value coded on OMPROUTE_IPV6_DEBUG_FILE. When either of these first files is full, the extensions are changed to 00N, where N is in the range of 1 to the number of files specified in the OMPROUTE_DEBUG_FILE_CONTROL environment variable (default 4). The current file is always the file named using the value coded on OMPROUTE_DEBUG_FILE/OMPROUTE_IPV6_DEBUG_FILE and the oldest file is the highest N value. This eliminates the danger of OMPROUTE filling the z/OS UNIX file system when tracing is active for a long time.

The size and number of debug files created can be controlled by the OMPROUTE_DEBUG_FILE_CONTROL environment variable. This allows you to adjust how much OMPROUTE trace data is saved. You tailor this parameter to your network complexity or available z/OS UNIX file system storage capacity. See z/OS Communications Server: IP Configuration Guide for details on this environment variable.

If the OMPROUTE CTRACE with option DEBUGTRC (or option ALL) is enabled, then output from OMPROUTE tracing and debugging is sent to the CTRACE facility. The OMPROUTE CTRACE facility can write trace records to an internal buffer or to an external writer. When the OMPROUTE CTRACE with option DEBUGTRC (or option ALL) is active, the normal debug output destinations are ignored. If the CTRACE is disabled, and a trace level is modified, then OMPROUTE once again follows the above rules for determining the debug output destination.

Sample OMPROUTE trace output

Figure 98 on page 778 is a sample OMPROUTE initialization and IPv4 routing protocol trace with descriptions for some of the trace entries:

```
1    EZZ7800I OMPROUTE starting
         EZZ7845I Established affinity with TCPCS8
     EZZ7817I Using defined OSPF protocol 89
     EZZ7838I Using configuration file: /u/user146/omproute/omproute.conf
2    EZZ7883I Processing interface from stack,address 9.169.100.18,
name CTC2,index 2,flags 451
EZZ7883I Processing interface from stack,address 9.67.100.8,
name CTC1,index 1,flags 451
EZZ8023I The RIP routing protocol is Enabled
2.5  EZZ8036I The IPv6 RIP routing protocol is Enabled
EZZ7937I The OSPF routing protocol is Enabled
EZZ8050I Updating BSD Route Parms for link CTC1, MTU 1024,
metric 1, subnet 255.255.255.0, destination 0.0.0.0
3    EZZ8057I Added network 9.67.100.0 to interface 9.67.100.8
                 on net 0 interface CTC1, table EZBMAIN
     EZZ7827I Adding stack route to 9.67.100.0, mask 255.255.255.0 via
            0.0.0.0, link CTC1, metric 1, type 1, table EZBMAIN
     EZZ8057I Added network 9.67.100.7 to interface 9.67.100.8 on net 0
            interface CTC1, table EZBMAIN
     EZZ7827I Adding stack route to 9.67.100.7, mask 255.255.255.255 via
            0.0.0.0, link CTC1, metric 1, type 129, table EZBMAIN
4    EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 0
                  interface CTC1
     EZZ7879I Joining multicast group 224.0.0.5 on interface 9.67.100.8
5    EZZ7913I State change, interface 9.67.100.8, new state 16,
                  event 1

⋮

     EZZ7875I No IPv4 Default Route Installed for table EZBMAIN
     EZZ81001 OMPROUTE subagent Starting
     EZZ7898I OMPROUTE Initialization Complete
     EZZ81011 OMPROUTE subagent Initialization Completed
     EZZ7908I Received packet type 1 from 9.167.100.13
6    EZZ8011I send request to address 9.67.100.7
     EZZ8015I sending packet to 9.67.100.7
     EZZ8011I send request to address 9.169.100.14
     EZZ8015I sending packet to 9.169.100.14
     EZZ8015I sending packet to 9.67.100.7
     EZZ8012I sending broadcast response to address 9.67.100.255 in 1
            packets with 1 routes
     EZZ8015I sending packet to 9.169.100.14
     EZZ8012I sending broadcast response to address 9.169.100.255 in 1
            packets with 1 routes
7    EZZ7908I Received packet type 1 from 9.67.100.7
     EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 0
            interface CTC1
8    EZZ7919I State change, neighbor 9.67.100.7, new state 4, event 1
9    EZZ7919I State change, neighbor 9.67.100.7, new state 8, event 3
     EZZ7934I Originating LS advertisement: typ 1 id 9.67.100.8
            org 9.67.100.8
10   EZZ7919I State change, neighbor 9.67.100.7, new state 16,
                  event 14
11   EZZ7910I Sending multicast, type 2, destination 224.0.0.5 net
                 0 interface CTC1
12   EZZ7908I Received packet type 2 from 9.67.100.7
13   EZZ7919I State change, IPv4neighbor 9.67.100.7, new state 32, event 5
14   EZZ7910I Sending multicast, type 3, destination 224.0.0.5 net 0
                  interface CTC1
     EZZ7908I Received packet type 2 from 9.67.100.7
15   EZZ7908I Received packet type 4 from 9.67.100.7
```

Figure 98. Sample OMPROUTE Trace Output

```
16   EZZ7928I from 9.67.100.7, new LS advertisement: typ 1 id
                  9.67.100.7 org 9.67.100.7
     EZZ7928I from 9.67.100.7, new LS advertisement: typ 1 id 9.67.100.8
            org 9.67.100.8
```

```
EZZ7927I from 9.67.100.7, self update: typ 1 id 9.67.100.8 org
       9.67.100.8
EZZ7928I from 9.67.100.7, new LS advertisement: typ 1 id
       9.167.100.13 org 9.100.13
EZZ7928I from 9.67.100.7, new LS advertisement: typ 5 id 9.67.100.0
       org 9.67.100.8
EZZ7927I from 9.67.100.7, self update: typ 5 id 9.67.100.0 org
       9.67.100.8
EZZ7928I from 9.67.100.7, new LS advertisement: typ 5 id 9.169.100.0
       org 9.67.100.8
EZZ7927I from 9.67.100.7, self update: typ 5 id 9.169.100.0 org
       9.67.100.8
EZZ7934I Originating LS advertisement: typ 1 id 9.67.100.8 org
       9.67.100.8
17 EZZ7910I Sending multicast, type 4, destination 224.0.0.5 net
            0 interface CTC1
EZZ7910I Sending multicast, type 3, destination 224.0.0.5 net 0
       interface CTC1
EZZ7908I Received packet type 4 from 9.67.100.7
EZZ7928I from 9.67.100.7, new LS advertisement: typ 5 id
       9.169.100.14 org 9.67.100.8
EZZ7927I from 9.67.100.7, self update: typ 5 id 9.169.100.14 org
       9.67.100.8
EZZ7910I Sending multicast, type 2, destination 224.0.0.5 net 0
       interface CTC1
EZZ7908I Received packet type 2 from 9.67.100.7
18 EZZ7919I State change, neighbor 9.67.100.7, new state 128,
            event 6
19 EZZ7908I Received packet type 5 from 9.67.100.7
20 EZZ7910I Sending multicast, type 5, destination 224.0.0.5
            net 0 interface CTC1
EZZ8015I sending packet to 9.169.100.14
EZZ8012I sending broadcast response to address 9.169.100.255 in
       1 packets with 1 routes
EZZ8015I sending packet to 9.67.100.7
EZZ8012I sending broadcast response to address 9.67.100.255 in
       1 packets with 1 routes
EZZ8015I sending packet to 9.169.100.14
EZZ8012I sending broadcast response to address 9.169.100.255 in
       1 packets with 1 routes
EZZ7908I Received packet type 4 from 9.67.100.7
EZZ7928I from 9.67.100.7, new LS advertisement: typ 1 id
       9.67.100.7 org 9.67.100.7
EZZ7910I Sending multicast, type 5, destination 224.0.0.5 net 0
       interface CTC1
EZZ7934I Originating LS advertisement: typ 5 id 9.169.100.14 org
       9.67.100.8
EZZ7934I Originating LS advertisement: typ 5 id 9.169.100.0 org
       9.67.100.8
EZZ7934I Originating LS advertisement: typ 5 id 9.67.100.0 org
       9.67.100.8
EZZ7910I Sending multicast, type 4, destination 224.0.0.5 net 0
       interface CTC1
21 EZZ7949I Dijkstra calculation performed, on 2 area(s), table EZBMAIN
EZZ7935I New OMPROUTE route to destination Net 9.67.100.7,
       type SPF cost 1, table EZBMAIN
EZZ7934I Originating LS advertisement: typ 3 id 9.67.100.7 org
       9.67.100.8
EZZ7908I Received packet type 5 from 9.67.100.7
EZZ7934I Originating LS advertisement: typ 1 id 9.67.100.8 org
       9.67.100.8
EZZ7910I Sending multicast, type 4, destination 224.0.0.5 net 0
       interface CTC1
EZZ7908I Received packet type 4 from 9.67.100.7
EZZ7928I from 9.67.100.7, new LS advertisement: typ 1 id
       9.167.100.13 org 9.167.100.13
EZZ7928I from 9.67.100.7, new LS advertisement: typ 4 id
```

```
       9.67.100.8 org 9.167.100.13
    EZZ7928I from 9.67.100.7, new LS advertisement: typ 3 id
       9.67.100.7 org 9.167.100.13
    EZZ7908I Received packet type 5 from 9.67.100.7
    EZZ7910I Sending multicast, type 5, destination 224.0.0.5 net 0
       interface CTC1
    EZZ7949I Dijkstra calculation performed, on 2 area(s), table EZBMAIN
22  EZZ7827I Adding stack route to 9.167.100.13, mask 255.255.
             255.255 via 9.67.100.7, link CTC1, metric 2, type 129, table EZBMAIN
    EZZ7935I New OMPROUTE route to destination Net 9.167.100.13,
       type SPF cost 2, table EZBMAIN
    EZZ7935I New OMPROUTE route to destination Net 9.67.100.8,
       type SPF cost 2, table EZBMAIN
    EZZ7913I State change, interface 9.67.100.8, new state 16, event 1
    EZZ7935I New OMPROUTE route to destination BR 9.167.100.13,
       type SPF cost 2, table EZBMAIN
    EZZ7827I Adding stack route to 9.167.100.17, mask 255.255.255.255
       via 9.67.100.7, link CTC1, metric 3, type 129, table EZBMAIN
    EZZ7935I New OMPROUTE route to destination Net 9.167.100.17,
       type SPF cost 3, table EZBMAIN
    EZZ7934I Originating LS advertisement: typ 3 id 9.167.100.13 org
       9.67.100.8
    EZZ7934I Originating LS advertisement: typ 3 id 9.67.100.8 org
       9.67.100.8
    EZZ7934I Originating LS advertisement: typ 3 id 9.167.100.17 org
       9.67.100.8
23  EZZ7909I Sending unicast type 1 dst 9.167.100.13
    EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 0
       interface CTC1
    EZZ7908I Received packet type 1 from 9.167.100.13
    EZZ7919I State change, neighbor 9.167.100.13, new state 4, event 1
    EZZ7919I State change, neighbor 9.167.100.13, new state 8, event 3
    EZZ7919I State change, neighbor 9.167.100.13, new state 16, event 14
    EZZ7909I Sending unicast type 2 dst 9.167.100.13
    EZZ7908I Received packet type 4 from 9.67.100.7
    EZZ7928I from 9.67.100.7, new LS advertisement: typ 4 id
       9.67.100.8 org 9.167.100.13
    EZZ7928I from 9.67.100.7, new LS advertisement: typ 3 id
       9.67.100.7 org 9.167.100.13
    EZZ7908I Received packet type 2 from 9.167.100.13
    EZZ7919I State change, neighbor 9.167.100.13, new state 32, event 5
    EZZ7909I Sending unicast type 2 dst 9.167.100.13
    EZZ7910I Sending multicast, type 5, destination 224.0.0.5 net 0
       interface CTC1
    EZZ7908I Received packet type 2 from 9.167.100.13
    EZZ7909I Sending unicast type 3 dst 9.167.100.13
    EZZ7908I Received packet type 4 from 9.167.100.13
    EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 1
       interface CTC2
    EZZ7928I from 9.167.100.13, new LS advertisement: typ 1 id
       9.67.100.8 org 9.67.100.8
    EZZ7927I from 9.167.100.13, self update: typ 1 id 9.67.100.8 org
       9.67.100.8
⋮
     EZZ7909I Sending unicast type 4 dst 9.167.100.13
    EZZ7919I State change, neighbor 9.167.100.13, new state 128, event 6
    EZZ7909I Sending unicast type 2 dst 9.167.100.13
    EZZ7934I Originating LS advertisement: typ 1 id 9.67.100.8 org
       9.67.100.8
    EZZ7910I Sending multicast, type 4, destination 224.0.0.5 net 0
       interface CTC1
    EZZ7933I Flushing advertisement: typ 3 id 9.67.100.7 org 9.167.100.13
    EZZ7933I Flushing advertisement: typ 4 id 9.67.100.8 org 9.167.100.13
    EZZ7909I Sending unicast type 5 dst 9.167.100.13
    EZZ8015I sending packet to 9.67.100.7
    EZZ8012I sending broadcast response to address 9.67.100.255 in
```

```
       1 packets with 1 routes
    EZZ7908I Received packet type 5 from 9.67.100.7
    EZZ7908I Received packet type 1 from 9.67.100.7
    EZZ8004I response received from host 9.67.100.7
    EZZ7908I Received packet type 5 from 9.167.100.13
    EZZ7949I Dijkstra calculation performed, on 2 area(s), table EZBMAIN
    EZZ8015I sending packet to 9.169.100.14
    EZZ8012I sending broadcast response to address 9.169.100.255 in
       1 packets with 1 routes
    EZZ7908I Received packet type 4 from 9.167.100.13
    EZZ7928I from 9.167.100.13, new LS advertisement: typ 1 id
       9.167.100.13 org 9.167.100.13
    EZZ7908I Received packet type 4 from 9.67.100.7
    EZZ7928I from 9.67.100.7, new LS advertisement: typ 1 id
       9.167.100.13 org 9.167.100.13
    EZZ7910I Sending multicast, type 5, destination 224.0.0.5 net 0
       interface CTC1
    EZZ7909I Sending unicast type 5 dst 9.167.100.13
    EZZ7934I Originating LS advertisement: typ 1 id 9.67.100.8 org
       9.67.100.8
    EZZ7909I Sending unicast type 4 dst 9.167.100.13
    EZZ7908I Received packet type 5 from 9.167.100.13
    EZZ8062I Subnet 9.0.0.0 defined, table EZBMAIN
    EZZ7949I Dijkstra calculation performed, on 2 area(s), table EZBMAIN
    EZZ7935I New OMPROUTE route to destination BR 9.167.100.13,
       type SPF cost 2, table EZBMAIN
24  EZZ7895I Processing DISPLAY command - OSPF,LIST,INTERFACES
    EZZ7809I    EZZ7833I INTERFACE CONFIGURATION
    EZZ7809I IP ADDRESS      AREA    COST  RTRNS  TRNSDLY PRI  HELLO   DEAD
    EZZ7809I 9.169.100.18    0.0.0.0    1     10      1     1     20    80
    EZZ7809I 9.67.100.8      2.2.2.2    1     10      1     1     20    80
    EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 0
       interface CTC1
    EZZ7908I Received packet type 1 from 9.167.100.13
    EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 1
       interface CTC2
    EZZ7909I Sending unicast type 1 dst 9.167.100.13
    EZZ7908I Received packet type 1 from 9.67.100.7
    EZZ8015I sending packet to 9.67.100.7
    EZZ8012I sending broadcast response to address 9.67.100.255 in
       1 packets with 1 routes
    EZZ8004I response received from host 9.67.100.7
    EZZ8015I sending packet to 9.169.100.14
    EZZ8012I sending broadcast response to address 9.169.100.255 in
       1 packets with 1 routes
25  EZZ7895I Processing MODIFY command - TRACE=2
25.5
    EZZ7895I Processing MODIFY command -TRACE6=2
    EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 0
       interface CTC1
26  EZZ7876I -- OSPF Packet Sent ------ Type: Hello
    EZZ7878I OSPF Version: 2            Packet Length: 48
    EZZ7878I Router ID: 9.67.100.8      Area: 2.2.2.2
    EZZ7878I Checksum: 1dcf             Authentication Type: 0
    EZZ7878I Hello_Interval: 20         Network mask: 255.255.255.0
    EZZ7878I Options: E
    EZZ7878I Router_Priority: 1         Dead_Router_Interval: 80
    EZZ7878I Backup DR: 0.0.0.0         Designated Router: 0.0.0.0
    EZZ7878I Neighbor: 9.67.100.7
    EZZ7877I -- OSPF Packet Received -- Type: Hello
    EZZ7878I OSPF Version: 2            Packet Length: 48
    EZZ7878I Router ID: 9.67.100.7      Area: 2.2.2.2
    EZZ7878I Checksum: 1dcf             Authentication Type: 0
    EZZ7878I Hello_Interval: 20         Network mask: 255.255.255.0
    EZZ7878I Options: E
    EZZ7878I Router_Priority: 1         Dead_Router_Interval: 80
    EZZ7878I Backup DR: 0.0.0.0         Designated Router: 0.0.0.0
```

```
EZZ7878I Neighbor: 9.67.100.8
EZZ7908I Received packet type 1 from 9.67.100.7
27          -- RIP Packet Received -- Type: Response (V1)
       Destination_Addr: 9.169.100.0     metric: 2
EZZ8004I response received from host 9.67.100.7
       -- RIP Packet Sent ------ Type: Response (V1)
       Destination_Addr: 9.169.100.0     metric: 1
EZZ8015I sending packet to 9.67.100.7
EZZ8012I sending broadcast response to address 9.67.100.255 in
       1 packets with 1 routes
28 EZZ7895I Processing MODIFY command - TRACE=1
EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 1
       interface CTC2
EZZ7909I Sending unicast type 1 dst 9.167.100.13
EZZ7908I Received packet type 1 from 9.67.100.7
EZZ8004I response received from host 9.67.100.7
EZZ8015I sending packet to 9.67.100.7
EZZ8004I response received from host 9.67.100.7
EZZ8015I sending packet to 9.67.100.7
EZZ8012I sending broadcast response to address 9.67.100.255 in 1
       packets with 1 routes
EZZ8015I sending packet to 9.169.100.14
EZZ8012I sending broadcast response to address 9.169.100.255 in 1
       packets with 1 routes
EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 0
       interface CTC1
```

⋮

```
EZZ7909I Sending unicast type 1 dst 9.167.100.13
EZZ7908I Received packet type 1 from 9.67.100.7
29 EZZ7862I Received update interface CTC1
30 EZZ8061I Deleted net 9.67.100.0 route via 9.67.100.8 net 0
             interface CTC1, table EZBMAIN
EZZ7864I Deleting all stack routes to 9.67.100.0, mask 255.255.255.0, table EZBMAIN
31 EZZ7919I State change, neighbor 9.67.100.7, new state 1, event 11
EZZ7879I Leaving multicast group 224.0.0.5 on interface 9.67.100.8
32 EZZ7913I State change, interface 9.67.100.8, new state 1, event 7
EZZ7934I Originating LS advertisement: typ 1 id 9.67.100.8 org
       9.67.100.8
EZZ7934I Originating LS advertisement: typ 1 id 9.67.100.8 org
       9.67.100.8
EZZ7909I Sending unicast type 4 dst 9.167.100.13
EZZ8015I sending packet to 9.169.100.14
EZZ8012I sending broadcast response to address 9.169.100.255 in
       1 packets with 1 routes
EZZ7934I Originating LS advertisement: typ 5 id 9.67.100.0 org
       9.67.100.8
EZZ7933I Flushing advertisement: typ 5 id 9.67.100.0 org 9.67.100.8
EZZ7949I Dijkstra calculation performed, on 1 area(s), table EZBMAIN
EZZ7801I Deleting stack route to 9.67.100.7, mask 255.255.255.255
       via 0.0.0.0, link CTC1, metric 1, type 129, table EZBMAIN
EZZ7935I New OMPROUTE route to destination Net 9.67.100.7,
       type SPIA cost 5, table EZBMAIN
EZZ7943I Destination Net 9.167.100.13 now unreachable, table EZBMAIN
EZZ7864I Deleting all stack routes to 9.167.100.13, mask
       255.255.255.255, table EZBMAIN
EZZ7935I New OMPROUTE route to destination Net 9.67.100.8,
       type SPIA cost 4, table EZBMAIN
EZZ7919I State change, neighbor 9.167.100.13, new state 1, event 11
EZZ7913I State change, interface 9.67.100.8, new state 1, event 7
EZZ7943I Destination BR 9.167.100.13 now unreachable, table EZBMAIN
EZZ7943I Destination Net 9.167.100.17 now unreachable, table EZBMAIN
EZZ7864I Deleting all stack routes to 9.167.100.17, mask
       255.255.255.255, table EZBMAIN
EZZ7934I Originating LS advertisement: typ 3 id 9.67.100.7
       org 9.67.100.8
EZZ7934I Originating LS advertisement: typ 3 id 9.67.100.8 org
```

```
       9.67.100.8
⋮
EZZ7933I Flushing advertisement: typ 3 id 9.167.100.17 org 9.67.100.8
EZZ7933I Flushing advertisement: typ 3 id 9.67.100.8 org 9.67.100.8
EZZ7933I Flushing advertisement: typ 3 id 9.167.100.13 org 9.67.100.8
EZZ7933I Flushing advertisement: typ 3 id 9.67.100.7 org 9.67.100.8
EZZ8015I sending packet to 9.169.100.14
EZZ8012I sending broadcast response to address 9.169.100.255 in 1
       packets with 1 routes
EZZ7949I Dijkstra calculation performed, on 1 area(s), table EZBMAIN
EZZ7943I Destination Net 9.67.100.7 now unreachable, table EZBMAIN
EZZ7943I Destination Net 9.67.100.8 now unreachable, table EZBMAIN
EZZ7943I Destination BR 9.167.100.13 now unreachable, table EZBMAIN
EZZ7934I Originating LS advertisement: typ 3 id 9.67.100.7 org
       9.67.100.8
EZZ7934I Originating LS advertisement: typ 3 id 9.67.100.8 org
       9.67.100.8
EZZ7933I Flushing advertisement: typ 3 id 9.67.100.8 org 9.67.100.8
EZZ7933I Flushing advertisement: typ 3 id 9.67.100.7 org 9.67.100.8
EZZ7910I Sending multicast, type 1, destination 224.0.0.5 net 1
       interface CTC2
EZZ7804I OMPROUTE exiting
```

Following are brief explanations of numbered items in the trace:

1 OMPROUTE initializing (trace level 1 was specified at startup: -t1).

2 OMPROUTE learns of TCP/IP stack IPv4 interfaces.

2.5 IPv6 tracing is in the file pointed to by the OMPROUTE_IPV6_DEBUG_FILE environment variable

3 Direct routes are added for each TCP/IP stack IPv4 interface.

4 OSPF Hello packet sent out OSPF interface.

5 OSPF Interface transitions to state “point-to-point.”

6 RIP Requests & Responses begin being sent out RIP interface.

7 OSPF Hello packet received from OSPF neighbor.

8 OSPF neighbor transitions to state “Init.”

9 OSPF neighbor transitions to state “2-Way.”

10 OSPF neighbor transitions to state “ExStart.”

11 OSPF Database Description packet sent out OSPF interface.

12 OSPF Database Description received from OSPF neighbor.

13 OSPF neighbor transitions to state “Exchange.”

14 OSPF Link State Request packet sent out OSPF interface.

15 OSPF Link State Update packet received from OSPF neighbor.

16 Link State Advertisements from received Update packet are processed.

17 OSPF Link State Update packet sent out OSPF interface.

18 OSPF neighbor transitions to state “Full.”

19 OSPF Link State Acknowledgment packet received from OSPF neighbor.

20 OSPF Link State Acknowledgment packet sent out OSPF interface.

21 OSPF Dijkstra calculation is performed.

22 Learned route is added to TCP/IP stack IPv4 route table.

23 Adjacency establishment begins with router at other end of OSPF Virtual Link.

24 Request received to display OSPF Interface configuration information.

25 Request received to change IPv4 tracing level to 2 (adds formatted packets).

25.5 Request received to change IPv6 tracing level to 2 (adds formatted packets to trace output in the file pointed to by the OMPROUTE_IPV6_DEBUG_FILE environment variable).

26 Formatted OSPF packet.

27 Formatted RIP packet.

28 Request received to change tracing level back to 1(-t1).

29 OMPROUTE learns of stopped TCP/IP IPv4 interface.

30 Routes over stopped interface are deleted.

31 Neighbor over stopped interface transitions to state "Down."

32 Stopped interface transitions to state "Down."

The following sample shows OMPROUTE IPv6 routing protocol trace with descriptions for some of the trace entries:

```
1  EZZ7977I Processing IPv6 interface from stack, address 1977::7,
name MPCPTPV67, index 16, flags 811, flags2 0
EZZ7977I Processing IPv6 interface from stack, address
fe80::542c:ed1e:1362:4d26, name MPCPTPV67, index 16, flags 811,
flags2 2
EZZ7977I Processing IPv6 interface from stack, address
7:7:7:7:7:7:7:7, name VIPA16, index 18, flags 4001, flags2 0
2  EZZ8057I Added network 1977::7 to interface
fe80::542c:ed1e:1362:4d26 on net 16 interface MPCPTPV67, table EZBMAIN
EZZ7879I Joining multicast group ff02::9 on interface MPCPTPV67
EZZ8057I Added network 7:7:7:7:7:7:7:7 to interface
7:7:7:7:7:7:7:7 on net 18 interface VIPA16, table EZBMAIN
EZZ8057I Added network 7:7:7:: to interface 7:7:7:7:7:7:7:7 on
net 18 interface VIPA16, table EZBMAIN
EZZ7827I Adding stack route to ::, prefixlen 0 via
fe80::7cb6:c5d5:6593:c076, link MPCPTPV67, metric 0, type 136, table EZBMAIN
3  EZZ8011I send request to address ff02::9
EZZ8015I sending packet to ff02::9
EZZ8021I sending IPv6RIP response to address ff02::9 from
fe80::542c:ed1e:1362:4d26 in 1 packets with 6 routes
4  EZZ8004I response received from host fe80::846e:70a6:8ca6:48b7
EZZ7827I Adding stack route to ::, prefixlen 0 via
fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 0, type 136, table EZBMAIN
EZZ7801I Deleting stack route to ::, prefixlen 0 via
fe80::7cb6:c5d5:6593:c076, link MPCPTPV67, metric 0, type 136, table EZBMAIN
EZZ8010I update route to net :: at metric 9 hops via router
fe80::846e:70a6:8ca6:48b7, table EZBMAIN
EZZ7806I Changing stack route to ::, prefixlen 0 via
fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 9, type 136, table EZBMAIN
EZZ8010I update route to net 1967::6 at metric 2 hops via router
fe80::846e:70a6:8ca6:48b7, table EZBMAIN
EZZ7827I Adding stack route to 1967::6, prefixlen 128 via
fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 2, type 1, table EZBMAIN
EZZ8010I update route to net 6:6:6:6:6:6:6:6 at metric 2 hops
via router fe80::846e:70a6:8ca6:48b7, table EZBMAIN
EZZ7827I Adding stack route to 6:6:6:6:6:6:6:6, prefixlen 128
via fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 2, type 1, table EZBMAIN
```

```
EZZ8010I update route to net 6:6:6:: at metric 2 hops via router
fe80::846e:70a6:8ca6:48b7, table EZBMAIN
EZZ7827I Adding stack route to 6:6:6::, prefixlen 48 via
fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 2, type 12, table EZBMAIN
EZZ8010I update route to net 1946::6 at metric 2 hops via router
fe80::846e:70a6:8ca6:48b7, table EZBMAIN
EZZ7827I Adding stack route to 1946::6, prefixlen 128 via
fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 2, type 1, table EZBMAIN
EZZ8010I update route to net 9::67:120:4 at metric 3 hops via
router fe80::846e:70a6:8ca6:48b7, table EZBMAIN
EZZ7827I Adding stack route to 9::67:120:4, prefixlen 128 via
fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 3, type 1, table EZBMAIN
EZZ8010I update route to net 1946::4 at metric 3 hops via router
fe80::846e:70a6:8ca6:48b7, table EZBMAIN
EZZ7827I Adding stack route to 1946::4, prefixlen 128 via
fe80::846e:70a6:8ca6:48b7, link MPCPTPV67, metric 3, type 1, table EZBMAIN
...
5  EZZ8015I sending packet to ff02::9
6   -- IPv6 RIP Packet Sent (MPCPTPV67) -- Type: Response
   Destination_Addr: ::
     Prefix Length: 0   metric: 16
   Destination_Addr: 9::67:120:3
     Prefix Length: 128   metric: 5
   Destination_Addr: 1977::7
     Prefix Length: 128   metric: 1
   Destination_Addr: 7:7:7:7:7:7:7:7
     Prefix Length: 128   metric: 1
   Destination_Addr: 7:7:7::
     Prefix Length: 48   metric: 1
   Destination_Addr: 9::67:120:7
     Prefix Length: 128   metric: 1
   Destination_Addr: 1967::
     Prefix Length: 16   metric: 1
   Destination_Addr: 1967::6
     Prefix Length: 128   metric: 16
   Destination_Addr: 6:6:6:6:6:6:6:6
     Prefix Length: 128   metric: 16
   Destination_Addr: 6:6:6::
     Prefix Length: 48   metric: 16
   Destination_Addr: 1946::6
     Prefix Length: 128   metric: 16
   Destination_Addr: 9::67:120:4
     Prefix Length: 128   metric: 16
   Destination_Addr: 1946::4
     Prefix Length: 128   metric: 16
   Destination_Addr: 1946::
     Prefix Length: 16   metric: 2
   Destination_Addr: 1111::
     Prefix Length: 16   metric: 16
   Destination_Addr: 50c9:c2d4::
     Prefix Length: 64   metric: 3
EZZ8021I sending IPv6RIP response to address ff02::9 from
fe80::542c:ed1e:1362:4d26 in 1 packets with 16 routes
...
EZZ8004I response received from host fe80::846e:70a6:8ca6:48b7
 -- IPv6 RIP Packet Received (MPCPTPV67) -- Type: Response
   Destination_Addr: ::
     Prefix Length: 0   metric: 10
   Destination_Addr: 1967::6
     Prefix Length: 128   metric: 1
   Destination_Addr: 1967::
     Prefix Length: 16   metric: 1
   Destination_Addr: 6:6:6:6:6:6:6:6
     Prefix Length: 128   metric: 1
   Destination_Addr: 6:6:6::
     Prefix Length: 48   metric: 1
   Destination_Addr: 1946::6
```

```
  Prefix Length: 128   metric: 1
Destination_Addr: 50c9:c2d4::
  Prefix Length: 64   metric: 16
Destination_Addr: 1111::
  Prefix Length: 16   metric: 16
Destination_Addr: 9::67:120:3
  Prefix Length: 128   metric: 16
Destination_Addr: 1977::7
  Prefix Length: 128   metric: 16
Destination_Addr: 7:7:7:7:7:7:7:7
  Prefix Length: 128   metric: 16
Destination_Addr: 7:7:7::
  Prefix Length: 48   metric: 16
Destination_Addr: 9::67:120:7
  Prefix Length: 128   metric: 16
Destination_Addr: 1976::
  Prefix Length: 16   metric: 2
Destination_Addr: f000::
  Prefix Length: 4   metric: 2
Destination_Addr: 9::67:120:4
  Prefix Length: 128   metric: 2
Destination_Addr: 1946::4
  Prefix Length: 128   metric: 2
```

Figure 99. Sample IPv6 OMPROUTE Trace Output

Following are brief explanations of numbered items in the trace:

1 OMPROUTE learns of TCP/IP stack IPv6 interface addresses. Each home address on an IPv6 interface is described separately; OMPROUTE uses the interface name to assign addresses to a specific interface.

2 Direct routes are added for each non-link-local TCP/IP stack IPv6 home address. When an interface's home address is needed in a message, its link-local address is used unless it is a VIPA that does not have a link-local address.

3 IPv6 RIP Requests and Responses begin being sent out IPv6 RIP interface. Note use of link-local address when interface is being identified by address only.

4 IPv6 RIP Response received and associated routes added to IPv6 route table. Source address is always link-local.

5 Request received to change IPv6 tracing level to 2 (adds formatted packets). The operator command to set the tracing level appears in the IPv4 trace, because **modify** commands run on the IPv4 thread.

6 Formatted IPv6 RIP packet.

TCP/IP services component trace for OMPROUTE

z/OS Communications Server provides Component Trace support for the OMPROUTE application. This section describes how to specify OMPROUTE trace and formatting options. For short descriptions of other tracing procedures, such as displaying trace status, see Chapter 5, "TCP/IP services traces and IPCS support," on page 45. Also, see "Commands to enable, disable, and display the status of the OMPROUTE CTRACE" on page 790.

For detailed descriptions, see the following information:

- z/OS MVS Diagnosis: Tools and Service Aids for information about Component Trace procedures
- z/OS MVS Initialization and Tuning Reference for information about the SYS1.PARMLIB member
- z/OS MVS System Commands for information about trace commands
- z/OS MVS Programming: Authorized Assembler Services Guide for information about procedures and return codes for CTRACE macros
- z/OS MVS IPCS Commands
- z/OS MVS IPCS User's Guide

Specifying trace options

You can specify Component Trace options at OMPROUTE initialization or after OMPROUTE has initialized.

Specifying options at initialization

A default minimum Component Trace is always started during OMPROUTE initialization. A parmlib member can be used to customize the parameters used to initialize the trace. The default OMPROUTE Component Trace parmlib member is the SYS1.PARMLIB member CTIORA00. The parmlib member name can be changed by use of the OMPROUTE_CTRACE_MEMBER environment variable.

Tip: Besides specifying the trace options, you can also change the OMPROUTE trace buffer size. The buffer size can be changed only at OMPROUTE initialization.

The maximum OMPROUTE trace buffer size is 100 MB.

Guideline: Use of a large internal CTRACE buffer or an external writer is recommended when using the DEBUGTRC option.

Requirement: The OMPROUTE REGION size in the OMPROUTE catalog procedure must be large enough to accommodate a large buffer size.

If the CTIORA00 member is not found when starting OMPROUTE, the following message is issued:

```
IEE5381 CTIORA00 MEMBER NOT FOUND in SYS1.PARMLIB
```

When this occurs, the OMPROUTE component trace is started with a buffer size of 1 MB and the MINIMUM tracing option.

The following figure shows the SYS1.PARMLIB member CTIORA00.

```
/*********************************************************************/
/*                                                                   */
/*  IBM Communications Server for z/OS                               */
/*  SMP/E Distribution Name: CTIORA00                                */
/*                                                                   */
/*  PART Name:  CTIORA00                                             */
/*                                                                   */
/*                                                                   */
/*  Copyright:                                                       */
/*               Licensed Materials - Property of IBM                */
/*               5694-A01                                            */
/*               (C) Copyright IBM Corp. 1998,2003                   */
/*                                                                   */
/*                                                                   */
/*  Status:      CSV1R5                                              */
/*                                                                   */
/*                                                                   */
```

```
/*  DESCRIPTION = This parmlib member causes component trace for     */
/*                the TCP/IP OMPROUTE application to be initialized  */
/*                with a trace buffer size of 1M                     */
/*                                                                   */
/*                This parmlib member only lists those TRACEOPTS     */
/*                values specific to OMPROUTE.  For a complete list  */
/*                of TRACEOPTS keywords and their values see         */
/*                z/OS MVS INITIALIZATION AND TUNING REFERENCE.      */
/*                                                                   */
/*                                                                   */
/* $MAC(CTIORA00),COMP(OSPF  ),PROD(TCPIP ):  Component Trace        */
/*                                            SYS1.PARMLIB member    */
/*                                                                   */
/*********************************************************************/
 TRACEOPTS
 /* ------------------------------------------------------------- */
 /*   Optionally start external writer in this file (use both     */
 /*   WTRSTART and WTR with same wtr_procedure)                   */
 /* ------------------------------------------------------------- */
 /*        WTRSTART(wtr_procedure)                                */
 /* ------------------------------------------------------------- */
 /*   ON OR OFF: PICK 1                                           */
 /* ------------------------------------------------------------- */
          ON
 /*       OFF                                                    */
 /* ------------------------------------------------------------- */
 /*   BUFSIZE: A VALUE IN RANGE 128K TO 100M                      */
 /*            CTRACE buffers reside in OMPROUTE Private storage  */
 /*            which is in the regions address space.             */
 /* ------------------------------------------------------------- */
          BUFSIZE(1M)
 /*       WTR(wtr_procedure)                                     */
 /* ------------------------------------------------------------- */
 /*   OPTIONS: NAMES OF FUNCTIONS TO BE TRACED, OR "ALL"          */
 /* ------------------------------------------------------------- */
 /*        OPTIONS(            */
 /*                'ALL     '  */
 /*               ,'MINIMUM '  */
 /*               ,'ROUTE   '  */
 /*               ,'PACKET  '  */
 /*               ,'OPACKET '  */
 /*               ,'RPACKET '  */
 /*               ,'IPACKET '  */
 /*               ,'SPACKET '  */
 /*               ,'DEBUGTRC'  */
 /*                )           */
```

Figure 100. SYS1.PARMLIB member CTIORA00

Table 75 describes the available trace options.

Table 75. OMPTRACE options

Trace event	Description
ALL	Select all types of records. Be aware that this option slows performance.
MINIMUM	Select OMPROUTE's minimum level of tracing. Specifying MINIMUM is the same as specifying ROUTE.

Table 75. OMPTRACE options (continued)

Trace event	Description
ROUTE	Select information exchange and routing updates between the OMPROUTE application and the z/OS TCP/IP Services stack.
PACKET	Select all inbound and outbound packet flows. This is the same as specifying OPACKET, RPACKET, and IPACKET.
RPACKET	Select inbound and outbound packet flows for the IPv4 RIP and IPv6 RIP protocols.
OPACKET	Select inbound and outbound packet flows for the IPv4 OSPF and IPv6 OSPF protocols.
IPACKET	Select inbound packets sent from z/OS TCP/IP with information regarding route or interface changes.
SPACKET	Trace inbound and outbound packets sent between the SNMP agent and the OMPROUTE subagent.
DEBUGTRC	Redirects IPv4 trace (-t), IPv4 debug (-d), IPv6 trace (-6t) and IPv6 debug (-6d) output to the CTRACE facility.

Guideline: Use of a large internal CTRACE buffer or an external writer is recommended when using the DEBUGTRC option.

Specifying options after initialization:
After OMPROUTE initialization, you must use the TRACE CT command to change the component trace options. Each time a new Component Trace is initiated, all prior trace options are turned OFF and the new options are put into effect.

You can specify the trace options with or without the PARMLIB member. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45.

Formatting OMPROUTE trace records

You can format component trace records using IPCS panels or a combination of the IPCS panels and the CTRACE command, either from a dump or from external-writer files. (For details, see Chapter 5, "TCP/IP services traces and IPCS support," on page 45.) Any combination of the following values can be entered as options to filter the CTRACE entries. The options must be entered using the format:

```
TYPE(option[,option]...)
```

- ROUTE
- OPACKET
- RPACKET
- IPACKET
- SPACKET
- DEBUGTRC

You cannot use the following as options when formatting OMPROUTE component traces:

- ALL
- MINIMUM
- PACKET

Commands to enable, disable, and display the status of the OMPROUTE CTRACE

Steps for enabling the CTRACE at OMPROUTE startup

Edit the CTIORA00 parmlib member as appropriate and start OMPROUTE with a trace level enabled.

Procedure

To enable the CTRACE at OMPROUTE startup, take the following steps:

Restriction: OMPROUTE must have READ access to the SYS1.PARMLIB data sets.

1. Edit CTIORA00 parmlib member (or the member specified in OMPROUTE_CTRACE_MEMBER environment variable) and specify TRACEOPTS ON, the desired buffer size by way of the BUFSIZE() parameter, and the desired CTRACE options. To direct the CTRACE to an external writer, also specify the name of the writer JCL procedure in the WTR() parameter. See the example CTIORA00 member.
2. Start OMPROUTE with a trace level enabled.

Steps for disabling the CTRACE at OMPROUTE startup

To disable the CTRACE at OMPROUTE startup, edit CTIORA00 or the member specified in OMPROUTE_CTRACE_MEMBER environment variable and specify TRACEOPTS OFF.

Steps for enabling the CTRACE after OMPROUTE is started

Enable the CTRACE after OMPROUTE is started.

Procedure

Perform the following steps:

1. Do one of the following:
   - Issue the following console commands to enable a CTRACE to an internal buffer:

     ```
     TRACE CT,ON,COMP=SYSTCPRT,SUB=(omproute_jobname)
     R xx,OPTIONS=(ctrace options),END
     ```

   - Issue the following console commands to enable a CTRACE to an external writer:

     ```
      TRACE CT,WTRSTART=writer_proc
      TRACE CT,ON,COMP=SYSTCPRT,SUB=(omproute_jobname)
      R xx,OPTIONS=(ctrace options),WTR=writer_proc,END
     ```

2. If DEBUGTRC or ALL is included in the CTRACE options, issue one of the following commands to modify the trace level:

   ```
   F,omproute_jobname,TRACE=x
   F,omproute_jobname,DEBUG=x
   F,omproute_jobname,TRACE6=x
   ```

   or

```
F,omproute_jobname,DEBUG6=xx
```

Requirement: This is required even if the OMPROUTE trace is already active.

Steps for disabling the CTRACE after OMPROUTE is started

Disable the CTRACE after OMPROUTE is started.

Procedure

Perform the following steps:

1. Issue the following console commands to disable a CTRACE to an internal buffer:

   ```
   TRACE CT,OFF,COMP=SYSTCPRT,SUB=(omproute_jobname)
   ```

   or

   Issue the following console commands to disable a CTRACE to an external writer:

   ```
   TRACE CT,OFF,COMP=SYSTCPRT,SUB=(omproute_jobname)
   TRACE CT,WTRSTOP=writer_proc
   ```

2. If DEBUGTRC or ALL is included in the CTRACE options, issue one of the following commands to modify the trace level:

   ```
   F,omproute_jobname,TRACE=x
   F,omproute_jobname,DEBUG=x
   F,omproute_jobname,TRACE6=x
   or
   F,omproute_jobname,DEBUG6=x
   ```

Step for displaying the CTRACE status

To display the CTRACE status, issue the following console command:

```
D TRACE,COMP=SYSTCPRT,SUB=(omproute_jobname)
```

Chapter 32. Diagnosing NCPROUTE problems

The NCPROUTE protocol provides a standardized interface, through which a server program on one host (NCPROUTE) can manage the routing tables and respond to SNMP route table requests for another program (Network Control Program).

This topic contains the following subtopics:

- "Definitions" on page 796
- "Diagnosing NCPROUTE problems" on page 797
- "NCPROUTE traces" on page 805

Figure 101 shows the NCPROUTE environment.

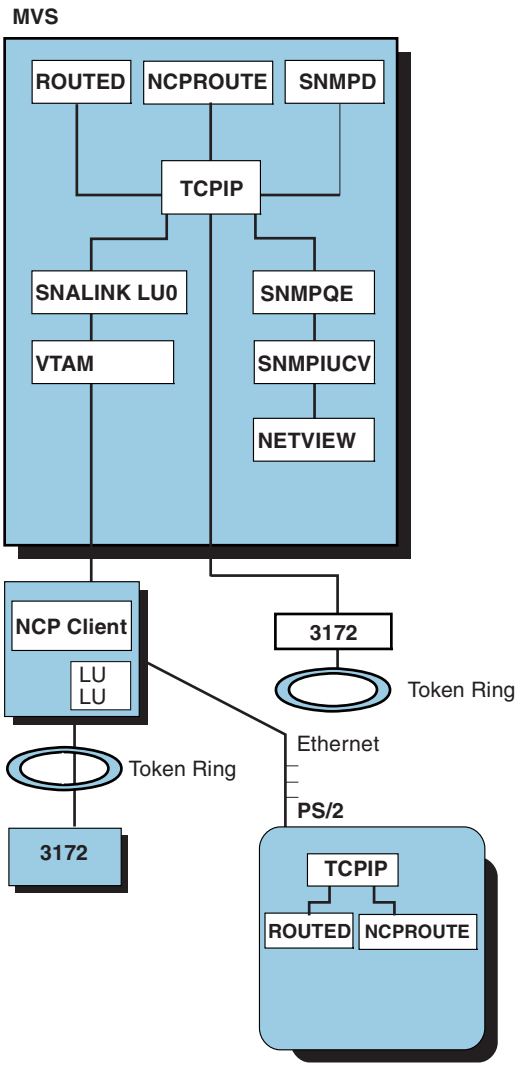

Figure 101. NCPROUTE environment

© Copyright IBM Corp. 2000, 2015

Prior to ACF/NCP V7R1, static route tables were used for routing IP datagrams over connected networks. However, the static routes had a drawback in that they were not able to respond to network topology changes. By implementing the RIP protocol between a host and NCP clients, the NCPROUTE server is able to provide dynamic IP routing for NCP clients. In effect, the NCP clients become active RIP routers in a TCP/IP network.

Multiple NCP units (374x family of communications controllers) can connect to the same NCPROUTE server on one host. This means that NCPROUTE can manage multiple routing tables for each NCP client. SNALINK is used as the connection vehicle to establish LU0 sessions between NCPROUTE and NCP clients. Each NCP client can have one or more LU0 sessions with NCPROUTE, provided that one session is used as primary and others as secondary for backup.

The NCPROUTE server reacts to network topology changes on behalf of NCP clients by maintaining each NCP client routing table, processing and generating RIP and SNMP datagrams, and performing error recovery procedures.

The NCPROUTE protocol is based on the exchange of protocol data units (PDUs).

The following list describes the eight types of PDUs:

Hello PDU
Sent from an NCP client to initiate a session with NCPROUTE.

Acknowledge PDU
Sent from NCPROUTE to acknowledge receipt of a Hello datagram. NCPROUTE is ready to manage the routing tables for an NCP client.

Status PDU
Sent from an NCP client to inform NCPROUTE of a status change with an interface. Interfaces can become inactive or active.

Delete Route Request PDU
Sent from NCPROUTE to request deletion of a route that is no longer known to the network from an NCP client routing table. This PDU can also be sent from an NCP client as a response informing NCPROUTE that the delete route request failed.

Add Route Request PDU
Sent from NCPROUTE to request addition of a route that is discovered by NCPROUTE to an NCP client routing table. This PDU can also be sent from an NCP client as a response informing NCPROUTE that the add route request failed.

Change Route Request PDU
Sent from NCPROUTE to request changing the value of a metric for a route currently active in an NCP client routing table.

Transport PDU
Sent from an NCP client to request NCPROUTE to retransmit RIP broadcasts sent from other routers and to process Simple Network Management Protocol (SNMP) requests sent from SNMP clients in the network. This PDU can also be sent from NCPROUTE as a response to retransmit RIP broadcasts or as a response to an SNMP query request. The Transport PDU contains encapsulated RIP and SNMP commands for additional processing.

Inactive Interface List PDU
Sent from an NCP client to inform NCPROUTE of currently inactive interfaces.

NCPROUTE uses the RIP messages for retransmitting of and responding to RIP updates and trace requests. A message might be unicasted, broadcasted, or multicasted, depending on the network interface capabilities in an NCP client.

There are four types of RIP messages that can be encapsulated in a Transport PDU. They are listed in Table 76

Table 76. Types of RIP messages

Message type	Description
Request	There are two types of Request messages: **REQUEST TO** Sent from NCP by NCPROUTE over a network interface to request routing table from one or more neighboring RIP routers. **REQUEST FROM** NCP received from one or more neighboring RIP routers as a request to transmit all or part of this NCP's routing table as supplied by NCPROUTE.
Response	There are two types of Response messages: **RESPONSE TO** Sent from NCP by NCPROUTE as a response to a request from a neighboring RIP router or sent from NCP by NCPROUTE for advertisements of RIP updates at periodic intervals over a network interface. The message contains all or part of this NCP's routing table as supplied by NCPROUTE. **RESPONSE FROM** NCP received from a neighboring RIP router as a response to a request from NCP by NCPROUTE or received from one or more neighboring RIP router for advertisements of RIP updates at periodic intervals. The message contains all or part of a neighboring router's routing table.
TraceOn	NCP received a request from a neighboring RIP router to enable the actions trace provided by NCPROUTE.
TraceOff	NCP received a request to a neighboring RIP router to disable tracing provided by NCPROUTE.

NCPROUTE communicates with the SNMP agent over the Distributed Program Interface (DPI) to process the SNMP commands. In this configuration, NCPROUTE becomes the SNMP subagent to provide values of registered MIB variables to the SNMP agent.

There are four types of SNMP commands that can be encapsulated within a Transport PDU:

Get Request
NCP received a request from a client to obtain one or more MIB variable values from an SNMP agent.

Get Next Request
NCP received a request from a client to obtain the next variable value in the MIB tree from an SNMP agent.

Get Response
Sent from NCP to its client as a response to an SNMP request.

Set Request
NCP received a request from a client to set or change the value of one or more MIB variables in an SNMP agent. This command is not supported by NCPROUTE.

See z/OS Communications Server: IP User's Guide and Commands for detailed information about the SNMP commands.

Table 77 describes the MIB variables registered for use by NCPROUTE:

Table 77. MIB variables registered for use by NCPROUTE

MIB variable	Description
ipRouteDest	Destination IP address of this route
ipRouteMetric1	Primary routing metric for this route
ipRouteMetric2	Alternative routing metric for this route
ipRouteMetric3	Another alternative routing metric for this route
ipRouteMetric4	Another alternative routing metric for this route
ipRouteNextHop	IP address of the next hop of this route
ipRouteType	Type of route
ipRouteProto	Routing mechanism by which this route was learned
ipRouteMask	Mask value for this route

See z/OS Communications Server: IP User's Guide and Commands for detailed information about the MIB variables.

Definitions

NCPROUTE must be defined correctly to both NCP and TCP/IP. UDP port 580 must be reserved for NCPROUTE. Routes to the NCP clients must be defined on the BEGINROUTES or the BSDROUTINGPARMS statement for NCPROUTE connectivity.

See z/OS Communications Server: IP Configuration Reference for detailed information about TCP/IP and NCPROUTE server definitions.

Internet interfaces (token ring and Ethernet) and NCST logical units for communication with the TCP/IP host must be defined for each NCP client through NCP generation.

Guideline: If you use SNMP to query routing information of NCP clients, the SNMP query engine and agent must be configured correctly. For NCPROUTE to communicate with the SNMP agent, the MVS host name or IP address and community name must be defined in the NCPROUTE profile, SEZAINST(NCPRPROF). The SNMP agent community name must also be defined in the *hlq*.PW.SRC data set for proper verification.

See z/OS Communications Server: IP Configuration Reference for detailed information about SNMP definitions.

Diagnosing NCPROUTE problems

Problems with NCPROUTE are generally reported under one of the following categories:

- Abends
- Connection problems
- Analyzing routing failures
- Incorrect output
- Session outages

Use the information provided in the following topics for problem determination and diagnosis of errors reported against NCPROUTE.

Abends

An abend during NCPROUTE processing should result in messages and error-related information being sent to the system console. A dump of the error is needed unless the symptoms match a known problem.

Documentation

Code a SYSUDUMP DD or SYSABEND DD statement in the cataloged procedure used to start NCPROUTE to ensure that a useful dump is obtained in event of an abend.

Analysis

See z/OS Problem Management or Chapter 3, “Diagnosing abends, loops, and hangs,” on page 21 for information about debugging dumps produced during NCPROUTE processing.

Connection problems

NCPROUTE connection problems are reported when NCPROUTE is unable to connect to TCP/IP, when NCP clients are unable to connect to the NCPROUTE server, when SNALINK LU0 is unable to connect between the NCPROUTE server and an NCP client, and when NCPROUTE is unable to connect to an SNMP agent. Generally, this type of problem is caused by an error in the configuration or definitions (either in VTAM, TCPIP, SNALINK, SNMP, NCP, or NCPROUTE).

Documentation

The following documentation should be available for initial diagnosis of NCPROUTE connection problems:

- Documentation for NCPROUTE connection failure
  - TCP/IP console log
  - *hlq*.PROFILE.TCPIP data set
  - TCPIP.DATA data set
  - NCPROUTE cataloged procedure
- Documentation for NCP client connection failure
  - NCPROUTE console log
  - NCPROUTE.PROFILE data set
  - NCP client network definitions data set (NCP generation)
- Documentation for SNALINK LU0 connection failure
  - SNALINK LU0 console log
  - VTAM APPL definitions for SNALINK LU0s
- Documentation for SNMP agent problems
  - SNMP console logs for SNMP agent and client
  - *hlq*.MIBDESC.DATA data set
  - *hlq*.PW.SRC data set
  - NetView log (if the SNMP client is on an MVS host)

More documentation that might be needed is discussed in the analysis section.

Analysis

Table 78 shows symptoms of connection problems and refers to the steps needed for initial diagnosis of the error.

Table 78. NCPROUTE connection problems

Connection Problem	Analysis Steps
NCP client connection failure	1, 2, 7, 8, 10, 13
NCPROUTE connection failure	1, 3, 5, 6, 7, 8, 10, 11, 13
SNALINK LU0 connection failure	1, 3, 7, 8, 10, 12
SNMP Agent connection failure	4, 9, 10, 13

"Steps for NCPROUTE connection problems" gives the diagnostic steps referred to in Table 78.

For TCP/IP configuration-related problems, see z/OS Communications Server: IP Configuration Reference for more information.

Steps for NCPROUTE connection problems:

Diagnose NCPROUTE connection problems.

Procedure

Perform the following steps:

1. For an NCP client, make sure that the internet interfaces (token ring and Ethernet) and NCST logical units for communication with the TCP/IP host are

defined correctly in an NCP generation. See the *ACF/NCP IP Router Planning and Installation Guide* for detailed information about NCP definitions.

a. Make sure that the NCPROUTE UDP port (UDPPORT keyword), coded on the IPOWNER statement in an NCP generation, matches the value defined in the .ETC.SERVICES data set. If it is not coded, the value used is the default UDP port 580.

b. Verify that the assigned port numbers and service names for NCPROUTE and the router are correct. Also make sure that the router service port 520 is defined in the .ETC.SERVICES data set. The NCP clients use this port as a destination port when broadcasting RIP packets to adjacent routers.

c. Make sure that NCST logical units for the SNALINK LU0s are defined correctly. A partner LU name (INTFACE keyword) for the SNALINK-NCST interface, coded on the LU statement in an NCST GROUP of an NCP generation, should match the LU name in a SNALINK LU0 DEVICE statement in the *hlq*.PROFILE.TCPIP data set.

d. Make sure that the remote LU name (REMLU keyword) for the SNALINK-NCST interface, coded on the LU statement in an NCP generation, matches the VTAM application name in the VTAM APPL definitions for SNALINK LU0s. For more information about SNALINK configuration and VTAM APPL definitions, see z/OS Communications Server: IP Configuration Guide.

e. Make sure that the NCST partner LU name (INTFACE keyword) for the SNALINK-NCST interface, coded on the IPOWNER and IPLOCAL statements in an NCP generation, matches the partner LU name in Step 1b.

f. Make sure that the IP address for the TCP/IP host (HOSTADDR keyword), coded on the IPOWNER statement in an NCP generation, matches the IP address for the SNALINK LU0 device name coded on the HOME statement in the *hlq*.PROFILE.TCPIP data set.

g. Make sure that the IP address for the SNALINK-NCST interface (LADDR keyword), coded on an IPLOCAL statement in an NCP generation, matches the IP address for the SNALINK LU0 link name coded on the BEGINROUTES statement in the *hlq*.PROFILE.TCPIP data set.

h. Make sure that the destination IP address for the SNALINK-NCST interface (P2PDEST keyword), coded on an IPLOCAL statement in an NCP generation, matches the IP address on the IPOWNER statement in Step 1e.

i. Make sure that IPLOCAL statements are defined for the directly-attached NCP internet interfaces (token ring and Ethernet) in an NCP generation. Verify the correctness of the IP addresses (LADDR keyword), metric values (METRIC keyword), protocol type (PROTOCOL keyword), and subnetwork masks (SNETMASK keyword).

2. Make sure that the appropriate NCP LOADLIB is used and that it contains correct network definitions. The NCP LOADLIB must be in the search list referred to by the //DD STEPLIB statement. Verify that a 374x communications controller to be in the session with NCPROUTE is loaded with the correct NCP load module.

3. Make sure that appropriate cataloged procedures for NCPROUTE (NCPROUT) and SNALINK (SNALPROC) are used, and verify the correctness of the data set references.

   - For the SNALINK cataloged procedure, make sure that the number of SNALINK sessions is large enough to allow multiple NCP sessions with NCPROUTE. This number is referred to by the MAXSESS keyword on the EXEC statement.

4. If using SNMP, make sure that the appropriate cataloged procedure for the SNMP agent (SNMPD) is used and verify the correctness of the data set references. Do likewise for an SNMP client (SNMPQE on MVS host).
5. Make sure that NCPROUTE is configured correctly in the *hlq*.PROFILE.TCPIP data set. The cataloged procedure name (NCPROUT) is referred to on AUTOLOG (optional), and PORT statements. UDP port 580 must be reserved for NCPROUTE.
6. Make sure that NCPROUTE is configured correctly in the ETC.SERVICES data set. See also Step 1a.
7. Make sure that SNALINK LU0 is configured correctly in the *hlq*.PROFILE.TCPIP data set. The SNALINK device name, LU name, and VTAM application address space name are referred to on the DEVICE statement. The SNALINK link name is referred to on the LINK, HOME, and BEGINROUTES statements. See also Steps 1b, 1c, 1e, and 1f.
   - If more than one NCP client is to be in session with NCPROUTE, repeat Step 7 to configure SNALINK LU0 for another session. TCP/IP definitions must be defined for each SNALINK LU0 session. If TCP/IP is currently running and another NCP client is to be added, another SNALINK LU0 can be configured using VARY TCPIP,,OBEYFILE commands. This allows TCP/IP to be reconfigured without having to shut down TCP/IP.
8. If you are using OMPROUTE, make sure that the routing parameters in the OMPROUTE configuration file (network interface definitions) and TCP/IP configuration (BSDROUTINGPARMS and BEGINROUTES statements) for the NCP clients are defined correctly. In addition, verify that direct and static routes to the NCP clients are defined correctly in TCP/IP BEGINROUTES statement.
9. If you are using SNMP, make sure that the SNMP agent is configured correctly in the *hlq*.PROFILE.TCPIP data set. If the SNMP client is on an MVS host, verify that the SNMP client address space is also configured. The cataloged procedure names, SNMPD, for the SNMP agent and client, are referred to on the AUTOLOG (optional), and PORT statements.
   - For the SNMP agent, make sure that the access authority information is defined correctly in the SEZAINST(EZBNRPRF) data set for the NCPROUTE profile, referenced in the NCPROUTE cataloged procedure.
10. If an NCP client is activated and ready to establish a session with NCPROUTE, make sure that the cataloged procedures for TCPIP, NCPROUTE, and SNALINK are all started. If you are using SNMP, make sure that the SNMP agent and client are started.
    a. Make sure that the SNALINK devices are started by the START statement in the *hlq*.PROFILE.TCPIP data set. The SNALINK devices can also be started by a VARY TCPIP,,OBEYFILE command or a VARY TCPIP,,START command.
    b. Make sure that VTAM command prompts at the system operator console are replied to; otherwise, a SNALINK session can be in a pending activation state.
    c. Make sure that the NCP client physical and logical lines for the internet interfaces (token ring and Ethernet) are active.
    d. Make sure that NCST lines are active for the SNALINK LU0 sessions.
    e. Make sure that VTAM cross-domain resource managers (CDRMs) are active in the MVS hosts.
11. For network connectivity problems, see Chapter 4, “Diagnosing network connectivity problems,” on page 25.

12. For SNMP problems, see Chapter 24, "Diagnosing Simple Network Management Protocol (SNMP) problems," on page 609.
13. For OMPROUTE problems, see Chapter 31, "Diagnosing OMPROUTE problems," on page 759. Ensure that the interface definitions in the BSDROUTINGPARMS statement in the *hlq*.PROFILE.TCPIP data set match the definitions in the corresponding interface definitions in the OMPROUTE configuration file. For more information about defining BSDROUTINGPARMS for NCPROUTE, see z/OS Communications Server: IP Configuration Reference.

Analyzing routing failures

Routing problems are the result of outages in a network and there are no alternative routing paths available for recovery. They can also be the result of incorrect configurations in the channel-attached and network-attached routers, as well as incorrect ARP entries, when applicable. PING and Traceroute commands to and from a z/OS host are useful diagnosis aids for problem determination.

In this topic, unless otherwise specified, the following command terms are used as described in Table 79.

Table 79. NCPROUTE command terms

Term	Description
PING	Refers to z/OS UNIX **oping**, TSO PING, and the **ping** commands used on other platforms.
Traceroute	Refers to z/OS UNIX **otracert**|**traceroute**, TSO TRACERTE, and the **traceroute** commands used on other platforms.
NETSTAT ROUTE	Refers to the z/OS UNIX **onetstat -r**, TSO NETSTAT ROUTE, and the **netstat route** commands used on other platforms.
NETSTAT GATE	Refers to the z/OS UNIX **onetstat -g** and TSO NETSTAT GATE commands. This command is available only on the z/OS platform.
NETSTAT ARP	Refers to the z/OS UNIX **onetstat -R ALL**, TSO NETSTAT ALL, and the **netstat arp** commands used on other platforms.

NCPROUTE routing failures are reported when a client is unable to get a positive response to a PING or Traceroute command for a remote host where there are NCPs acting as RIP servers along the routing paths.

Documentation

The following documentation should be available for initial diagnosis of routing failures:

- NCPROUTE console log
- TCP/IP console log
- *hlq*.PROFILE.TCPIP data set
- NCP client network definitions data set (NCP generation)
- Output from MODIFY NCPROUTE,TABLES command for a display of internal tables representing an NCP client.
- Outputs from PING and Traceroute commands.

Analysis

Table 80 shows symptoms of PING failures and refers to the steps needed for initial diagnosis of the error.

Table 80. NCPROUTE routing failures

Routing Failure	Analysis Steps
Incorrect response	1, 2, 3, 4, 5, 6, 7, 8
Timeouts	2, 9

"Steps for analyzing routing failures" gives the diagnostic steps referred to in Table 80.

Steps for analyzing routing failures:

Analyze routing failures.

About this task

This topic gives the diagnostic steps referred to in Table 80.

Guideline: Because an NCP client cannot respond to Traceroute commands, you can use the PING command to diagnose routing failures. However, a Traceroute command can be used to locate a suspect router along the routing path to a remote host beyond the NCP client. In these steps, the PING command is used for diagnosis.

Procedure

Perform the following steps:

1. Make sure the PING command contains a valid destination IP address for the remote host.
2. Make sure that a 374x communications controller acting as a RIP server involved in the PING transaction is active and is running with a correct level of NCP LOADLIB. Verify that correct network definitions are defined in the NCP generation and that the NCP client is in session with NCPROUTE.
3. If the PING command was issued from a remote host, issue the NETSTAT ROUTE, NETSTAT GATE (if host is z/OS), and NETSTAT ARP commands from there for its routing and ARP table information.
   a. If the local host is running with OMPROUTE, verify the routing configuration for routes and networks as defined in the OMPROUTE configuration file (network interface definitions) and TCP/IP configuration (BSDROUTINGPARMS and BEGINROUTES statements). To ensure NCP connectivity with the NCP clients, verify that direct and static routes to the clients are defined correctly.
   b. If there are any problems with the routes and networks, see "Using the Netstat command" on page 39.
4. If the remote host is running with OMPROUTE, verify its routing configuration for routes and networks as defined in its OMPROUTE configuration file and TCP/IP configuration. See Step 3a for configuration information.
   a. For routers or hosts running on platforms other than z/OS, see their documentation for more information about correcting routing problems. Also, see these documentations for NETSTAT commands to display the routing and ARP tables for problem determination.

5. If there are no problems with the routes or networks, check for broken or poorly connected cables between the client and the remote host. This includes checking the IP interfaces (token ring and Ethernet) on the 374x communications controller.
6. Make sure there is a channel connection between the 374x communications controller and the MVS host. A channel connection can be interrupted by an Automatic Network Shutdown (ANS) situation. ANS can occur when the system operator puts the MVS console into CP mode. In this case, the system operator needs to return to MVS from CP to recover from ANS.
7. For more information about diagnosing network connectivity problems, see Chapter 4, "Diagnosing network connectivity problems," on page 25.
8. For more information about diagnosing PING problems, see "Using the Ping command" on page 34.
9. For more information about diagnosing PING timeouts, see "Correcting timeout problems" on page 38.

Incorrect output

Problems with incorrect output are reported when the data sent to the client is in an unexpected form (for example, incorrect TCP/IP output, incorrect SNALINK LU0 output, invalid RIP commands, incorrect RIP broadcasting information, incorrect routing-table updates, truncated packets, or incorrect SNMP agent or client output).

Documentation

The following documentation should be available for initial diagnosis of incorrect output:

- NCPROUTE cataloged procedure
- Documentation for NCPROUTE incorrect output
  - NCPROUTE console log
  - NCPROUTE.PROFILE data set
  - NCP client network definitions data set (NCP generation)
  - Output from MODIFY NCPROUTE,TABLES command for a display of internal tables (routes, interfaces, and filters) in NCPROUTE used for an NCP client.
- Documentation for TCP/IP incorrect output
  - TCP/IP console log
  - *hlq*.TCPIP. PROFILE data set
  - TCPIP.DATA data set
- Documentation for SNMP agent incorrect output
  - SNMP console logs for SNMP agent and client
  - *hlq*.MIBDESC.DATA data set
  - *hlq*
  - *hlq*.PW.SRC data set
  - NetView log (if SNMP client is on an MVS host)

Analysis

Table 81 on page 804 shows types of incorrect output and refers to the steps needed for initial diagnosis of the error.

Table 81. NCPROUTE incorrect output

Incorrect output	Analysis steps
TCP/IP incorrect output	1
SNALINK LU0 incorrect output	2
NCPROUTE incorrect output	3
SNMP agent or client incorrect output	4

Steps for diagnosing incorrect output:

Diagnose incorrect output.

About this task

This topic provides the diagnostic steps referred to in Table 81.

Procedure

Perform the following steps:

1. If the TCP/IP console shows a message, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD) and follow the directions for system programmer response for the message.
   a. Information in the TCP/IP console log should contain a detailed description of the error.
   b. In the event of TCP/IP loops or hangs, see Chapter 3, "Diagnosing abends, loops, and hangs," on page 21.
2. If the SNALINK LU0 console shows a SNALINK error, see the explanation of the corresponding error message as described in z/OS Communications Server: IP Messages Volume 1 (EZA) or z/OS Communications Server: SNA Messages. For more information about diagnosing SNALINK LU0 session outages, see Chapter 19, "Diagnosing SNALINK LU0 problems," on page 557.
3. If the NCPROUTE console shows a message, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD) and follow the directions for system programmer response for the message.
4. If the SNMP agent or client console shows a message, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD) and follow the directions for system programmer response for the message.
5. For more information about diagnosing SNMP problems, see Chapter 24, "Diagnosing Simple Network Management Protocol (SNMP) problems," on page 609.

Session outages

Session outages are reported as an unexpected termination of the TCP/IP connection, the SNALINK LU0 task, the NCPROUTE-to-NCP client session, or the NCPROUTE-to-SNMP agent connection. A session that was disconnected or ended results in NCPROUTE being returned to the initial state of waiting for Hello PDUs and SNMP requests from an NCP client.

Documentation

The following documentation should be available for initial diagnosis of session outages:

- Documentation for TCP/IP session outage

  - TCP/IP console log
- Documentation for SNALINK LU0 session outage
  - SNALINK LU0 console log
  - VTAM console log
- Documentation for NCPROUTE-to-NCP client session outage
  - NCPROUTE cataloged procedure
  - NCPROUTE console log
  - NCP client network definitions data set (NCP generation)
- Documentation for NCPROUTE-to-SNMP agent session outage
  - SNMP console log for SNMP agent
  - Tivoli NetView for z/OS log (if the SNMP client is on the MVS host)

Analysis

Table 82 shows symptoms of session outages and refers to the steps needed for initial diagnosis of the error.

Table 82. Symptoms of session outages

If this is the outage type	Then perform these steps
TCP/IP session outage	If the TCP/IP console shows a TCP/IP error message, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD) and follow the directions for system programmer response for the message. If TCP/IP abended, see Chapter 3, "Diagnosing abends, loops, and hangs," on page 21.
SNALINK LU0 session outage	If the SNALINK LU0 console shows a SNALINK error, see the explanation of the corresponding error message as described in z/OS Communications Server: IP Messages Volume 1 (EZA) or z/OS Communications Server: SNA Messages. For more information about diagnosing SNALINK LU0 session outages, see Chapter 19, "Diagnosing SNALINK LU0 problems," on page 557.
NCPROUTE-to-NCP client session outage	If the NCPROUTE console shows an NCPROUTE error message, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD) and follow the directions for system programmer response for the message.
NCPROUTE-to-SNMP agent session outage	If the SNMP agent console shows an SNMP error message, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD) and follow the directions for system programmer response for the message. For more information about diagnosing SNMP problems, see Chapter 24, "Diagnosing Simple Network Management Protocol (SNMP) problems," on page 609.

You can now perform the steps for the decision you have made.

NCPROUTE traces

There are many TCP/IP traces that can be useful in identifying the cause of NCPROUTE problems. This topic describes the NCPROUTE traces.

Guideline: NCPROUTE trace output is sent to the location specified by the SYSPRINT DD statement in the NCPROUTE cataloged procedure.

Activating NCPROUTE global traces

The NCPROUTE global traces are all controlled by parameters on `PARMS=` in the PROC statement of the NCPROUTE cataloged procedure. (*Global tracing* means that all NCP clients are traced.)

For example:

```
//NCPROUT PROC MODULE=NCPROUTE,PARMS='/-t -t'
```

Tip: These parameters are also valid when starting the NCPROUTE server with the START command.

The NCPROUTE parameters that control global tracing are:

-t Activates global tracing of actions for all NCP clients.

-t -t Activates global tracing of packets for all NCP clients. NCPROUTE tracing can be started and stopped using the MODIFY command. For more information, see z/OS Communications Server: IP System Administrator's Commands.

-tq Deactivates tracing at all levels. This parameter suppresses tracing for all NCP clients and overrides the trace settings on the GATEWAY statements in the NCPROUTE GATEWAYS data set.

-dp Activates global tracing of data packets coming in and out of NCPROUTE. The data is displayed in data format.

-dq Deactivates global tracing of data packets coming in and out of NCPROUTE.

Restrictions:

- A slash (/) must precede the first parameter.
- Each parameter must be separated by a blank.
- Mixed case is allowed for the parameters.
- The parameters for the NCPROUTE procedure are case-sensitive.
- There are no third- or fourth-level global tracing options like those on the GATEWAY statements in the NCPROUTE GATEWAYS data set. The system uses the higher of the two settings for a specific NCP client.
- The data packets trace option is not available for selective tracing.

The parameters described her include only those that activate tracing. See the z/OS Communications Server: IP Configuration Reference for more information about all of the NCPROUTE parameters.

Activating NCPROUTE selective traces

The NCPROUTE selective traces are all activated as trace options specified in the OPTIONS statement for an NCP client in the NCPROUTE GATEWAYS data set. Selective tracing means a different trace level can be specified for each NCP client. To assist in problem isolation, a particular NCP client can be selected for tracing.

The keyword on the OPTIONS statement that controls selective tracing for an NCP client is `trace.level`. The value that follows this keyword indicates the trace level to be used.

Value	Meaning
0	Does not activate any traces.
1	Activates tracing of actions by the NCPROUTE server.
2	Activates tracing of all packets sent or received.
3	Activates tracing of actions, packets sent or received, and packet history. Circular trace buffers are used for each interface to record the history of all packets traced. This history is included in the trace output whenever an interface becomes inactive.
4	Activates tracing of actions, packets sent or received, packet history, and packet contents. The RIP network routing information is included in the trace output.

Restriction: The selective traces must be defined prior to activation of an NCP client or prior to starting the NCPROUTE cataloged procedure.

See the z/OS Communications Server: IP Configuration Reference for more information about the GATEWAYS data set and the GATEWAY and OPTIONS statements.

For example, the following command would activate tracing of actions, packets sent or received, packet history, and packet contents:

```
options trace.level 4
```

NCPROUTE trace example and explanation

Figure 102 on page 808 shows an example of an NCPROUTE trace with actions, packets, history, and contents traced. The trace was generated with trace level 4 specified in the OPTIONS statement and `PARMS='/-t -dp'` in the PROC statement of the NCPROUTE cataloged procedure.

The trace level column does not appear in the actual trace. It was added to the example to indicate the levels of the trace for which the line is generated. For example, including: `trace.level 3` on the options statement NCP client GATEWAYS data set would result in a level 3 trace, and all of the lines indicated as trace level 1, 2, or 3 would be generated in the trace output. Lines indicated as trace level d are generated if the -dp parameter is specified.

```
Trace
level
0      1 15:29:48 EZB3826I Port 580 assigned to ncprout
0        15:29:49 EZB3885I Input parameter(s): -t -t -dp
1        15:29:49 EZB4159I Global tracing actions started
2        15:29:49 EZB4160I Global tracing packets started
0        15:29:49 EZB3834I ******************************************************
0      2 15:29:49 EZB4196I * Opening NCPROUTE profile dataset (DD:NCPRPROF)
0        15:29:49 EZB3834I ******************************************************
0      3 15:29:49 EZB4055I ** Attempting to (re)start SNMP connection
0        15:29:49 EZB4059I Connecting to agent 9.67.116.66 on DPI port 1141
0        15:29:49 EZB4062I SNMP DPI connection established
0        15:29:50 EZB4064I 1.3.6.1.4.1.2.6.17. registered with SNMP agent0
1        15:29:50 EZB3829I Waiting for incoming packets
d      4 ============== Received datagram from NCP client (length=32)
d        0000    0100 0000 c1f0 f4d5
d        0008    f7f1 f1d7 f0f5 61f0
d        0010    f661 f9f4 40f1 f07a
d        0018    f1f0 7af4 f200 0000
d        0020(32)
0        15:29:51 EZB3834I ******************************************************
0        15:29:51 EZB3876I * Hello from new client 9.67.116.65
0        15:29:51 EZB3877I * RIT dataset name: A04N711P
0        15:29:51 EZB3878I * RIT ID:  05/06/94 10:10:42
0        15:29:51 EZB3834I ******************************************************
0        15:29:51 EZB3867I Acknowledge to 9.67.116.65: Hello Received
0        15:29:51 EZB3999I Establishing session with client 9.67.116.65
0        15:29:51 EZB3868I Acknowledge to 9.67.116.65: RIT Loaded OK
0        15:29:51 EZB4166I Session with client 9.67.116.65 started
1        15:29:51 EZB3829I Waiting for incoming packets
1        ============== Received datagram from NCP client (length=8)
d        0000    0800 0000 0a44 005c
d        0008(8)
0        15:29:51 EZB3834I ******************************************************
0      5 15:29:51 EZB3898I * Recv:  Inactive Interface List from 9.67.116.65
0                           * 1 interface(s) found:
0        15:29:51 EZB3899I *   10.68.0.92 - TR92
0        15:29:51 EZB3834I ******************************************************
0        15:29:51 EZB3834I ******************************************************
0      6 15:29:51 EZB3956I * Processing interface NCSTALU1
0        15:29:51 EZB3834I ******************************************************
0        15:29:51 EZB3959I Point-to-point interface, using dstaddr
0        15:29:51 EZB3962I Adding (sub)network address for interface
1        15:29:51 EZB3912I ifwithnet:  compare with NCSTALU1
1        15:29:51 EZB3915I netmatch 9.67.116.65 and 9.67.116.65
1        15:29:51 EZB4029I Tue Jun 28 15:29:51:
1      7 15:29:52 EZB4030I ADD destination 9.67.116.66, router 9.67.116.66, metric 1
1                           flags UP|HOST state INTERFACE|CHANGED|INTERNAL|PERM|SUBNET timer 0
```

Figure 102. NCPROUTE trace

```
0        15:29:52 EZB3834I ******************************************************
0        15:29:52 EZB3956I * Processing interface TR88
0        15:29:52 EZB3834I ******************************************************
0        15:29:52 EZB3960I This interface is not point-to-point
0        15:29:52 EZB3962I Adding (sub)network address for interface
1        15:29:52 EZB3912I ifwithnet:  compare with NCSTALU1
1        15:29:52 EZB3912I ifwithnet:  compare with TR88
1        15:29:52 EZB3915I netmatch 10.68.0.88 and 10.68.0.88
1        15:29:52 EZB4030I ADD destination 10.0.0.0, router 10.68.0.88, metric 1
1                           flags UP state INTERFACE|CHANGED|INTERNAL|SUBNET|PERM timer 0
1        15:29:52 EZB3912I ifwithnet:  compare with NCSTALU1
1        15:29:52 EZB3912I ifwithnet:  compare with TR88
1        15:29:52 EZB3915I netmatch 10.68.0.88 and 10.68.0.88
1        15:29:52 EZB4030I ADD destination 10.68.0.0, router 10.68.0.0, metric 1
1                           flags UP state INTERFACE|CHANGED|SUBNET|PERM timer 0
0        15:29:52 EZB3834I ******************************************************
```

```
0        15:29:52 EZB3956I * Processing interface TR92
0        15:29:52 EZB3834I ******************************************************
0        15:29:52 EZB3960I This interface is not point-to-point
0        15:29:52 EZB3948I Interface TR92 not up
0        15:29:52 EZB3834I ******************************************************
0      8 15:29:52 EZB3973I * Opening GATEWAYS dataset for client 9.67.116.65
0                         *   'TCPCS.NCPROUTE.GATEWAYS(A04N711P)'
0        15:29:52 EZB3834I ******************************************************
0        15:29:52 EZB3968I Start of GATEWAYS processing:
0        15:29:52 EZB4195I Option(s): trace.level 4 supply on default.router no
1        15:29:52 EZB4015I Client tracing actions started
2        15:29:52 EZB4016I Client tracing packets started
3        15:29:52 EZB4017I Client tracing history started
4        15:29:52 EZB4018I Client tracing packet contents started
0        15:29:52 EZB4198I (no etc.gateway definitions)
0        15:29:52 EZB4150I End of GATEWAYS processing
1        15:29:52 EZB3829I Waiting for incoming packets
d      9 ============== Received datagram from NCP client (length=80)
d     10 0000    0700 0000 9200 004a
         0008    4500 0048 09c0 0000
d        0010    3c11 79dc 0943 7442
d        0018    0943 7441 0208 0208
d        0020    0034 079e 0201 0000
d        0028    0002 0000 0943 7441
d        0030    0000 0000 0000 0000
d        0038    0000 0001 0002 0000
d        0040    0943 7000 0000 0000
d        0048    0000 0000 0000 0001
d        0050(80)
d        ============== Transport PDU header (length=8)
d        0000    0700 0000 9200 004a
d        0008(8)
d        ============== IP header (length=20)
d        0000    4500 0048 09c0 0000
d        0008    3c11 79dc 0943 7442
d        0010    0943 7441 8002 c12c
d        0018(24)
d        ============== UDP header (length=8)
d        0000    0208 0208 0034 079e
d        0008(8)
d        ============== UDP data (length=44)
d        0000    0201 0000 0002 0000
d        0008    0943 7441 0000 0000
d        0010    0000 0000 0000 0001
d        0018    0002 0000 0943 7000
d        0020    0000 0000 0000 0000
d        0028    0000 0001 0001 6e68
d        0030(48)
1     11 15:30:04 EZB3894I Transport from 9.67.116.65:  44 bytes of RIP data
2        15:30:04 EZB4045I     RESPONSE from 9.67.116.66 -> 520:
d     12 ============== RIP net info (length=20)
d        0000    0002 0000 0943 7441
d        0008    0000 0000 0000 0000
d        0010    0000 0001 8002 c12c
d        0018(24)
4        15:30:04 EZB4049I         destination 9.67.116.65 metric 1
d        ============== RIP net info (length=20)
d        0000    0002 0000 0943 7000
d        0008    0000 0000 0000 0000
d        0010    0000 0001 8002 c12c
4        0018(24)
1        15:30:04 EZB4049I         destination 9.67.112.0 metric 1
1        15:30:04 EZB4029I Tue Jun 28 15:30:04:
1     13 15:30:04 EZB4030I ADD destination 9.67.112.0, router 9.67.116.66, metric 2
1                         flags UP|GATEWAY state CHANGED|SUBNET timer 0
1     14 15:30:04 EZB3855I NCP_Add out to 9.67.116.65
1                         Route to: 9.67.112.0 via interface 9.67.116.65 to 9.67.116.66
```

```
1                        Metric: 2, Type Subnet
1       15:30:04 EZB3829I Waiting for incoming packets
1    15 15:30:20 EZB4011I client 9.67.116.65:  30 second timer expired (broadcast)
1       15:30:20 EZB3951I client 9.67.116.65:  supply 9.67.116.66 -> 0 via NCSTALU1
4    16 15:30:20 EZB4045I     RESPONSE to 9.67.116.66 -> 0:
d       =============== RIP net info (length=20)
d       0000    0002 0000 0943 7442
d       0008    0000 0000 0000 0000
d       0010    0000 0001 8002 c12c
d       0018(24)
4       15:30:20 EZB4049I         destination 9.67.116.66 metric 1
d       =============== RIP net info (length=20)
d       0000    0002 0000 0a00 0000
d       0008    0000 0000 0000 0000
d       0010    0000 0001 8002 c12c
d       0018(24)
4       15:30:20 EZB4049I         destination 10.0.0.0 metric 1
d       =============== RIP net info (length=20)
d       0000    0002 0000 0943 7000
d       0008    0000 0000 0000 0000
d       0010    0000 0002 8002 c12c
d       0018(24)
4       15:30:20 EZB4049I         destination 9.67.112.0 metric 2
d       =============== UDP data (length=64)
d       0000    0201 0000 0002 0000
d       0008    0943 7442 0000 0000
d       0010    0000 0000 0000 0001
d       0018    0002 0000 0a00 0000
d       0020    0000 0000 0000 0000
d       0028    0000 0001 0002 0000
d       0030    0943 7000 0000 0000
d       0038    0000 0000 0000 0002
d       0040(64)
d       =============== UDP header (length=8)
d       0000    0208 0208 0048 fd70
d       0008(8)
d       =============== IP header (length=20)
d       0000    4500 005c 0000 0000
d       0008    0411 bb88 0943 7441
d       0010    0943 7442 8002 c12c
d       0018(24)
d       =============== Transport PDU header (length=8)
d       0000    0700 0000 0943 7441
d       0008(8)
d       =============== Sending Transport PDU to NCP client (length=100)
d       0000    0700 0000 0943 7441
d       0008    4500 005c 0000 0000
d       0010    0411 bb88 0943 7441
d       0018    0943 7442 0208 0208
d       0020    0048 fd70 0201 0000
d       0028    0002 0000 0943 7442
d       0030    0000 0000 0000 0000
d       0038    0000 0001 0002 0000
d       0040    0a00 0000 0000 0000
d       0048    0000 0000 0000 0001
d       0050    0002 0000 0943 7000
d       0058    0000 0000 0000 0000
d       0060    0000 0002 0000 0001
d       0068(104)
.
.
.
1    17 15:30:20 EZB3948I Interface TR92 not up
1       15:30:20 EZB3829I Waiting for incoming packets
                         .
1       15:30:20 EZB3894I Transport from 9.67.116.65:  64 bytes of RIP data
2       15:30:20 EZB4045I     RESPONSE from 10.68.0.88 -> 520:
                         .
4       15:30:20 EZB4049I         destination 9.67.116.66 metric 1
```

```
                          .
4       15:30:20 EZB4049I         destination 10.68.0.0 metric 1
                          .
4       15:30:20 EZB4049I         destination 9.67.112.0 metric 2
1       15:30:20 EZB3829I Waiting for incoming packets
                          .
                          .
                          .
1        15:30:34 EZB3829I Waiting for incoming packets
1        15:30:50 EZB4011I client 9.67.116.65:  30 second timer expired (broadcast)
1        15:30:50 EZB3951I client 9.67.116.65:  supply 9.67.116.66 -> 0 via NCSTALU1
2        15:30:50 EZB4045I     RESPONSE to 9.67.116.66 -> 0:
                          .
4        15:30:50 EZB4049I         destination 9.67.116.66 metric 1
                          .
4        15:30:50 EZB4049I         destination 10.0.0.0 metric 1
                          .
4        15:30:50 EZB4049I         destination 9.67.112.0 metric 2
                          .
1        15:30:50 EZB3951I client 9.67.116.65:  supply 10.68.15.255 -> 0 via TR88
2        15:30:50 EZB4045I     RESPONSE to 10.68.15.255 -> 0:
                          .
4        15:30:50 EZB4049I         destination 9.67.116.66 metric 1
                          .
4        15:30:50 EZB4049I         destination 10.68.0.0 metric 1
                          .
4        15:30:50 EZB4049I         destination 9.67.112.0 metric 2
:
1        15:32:35 EZB3894I Transport from 9.67.116.65:  64 bytes of RIP data
2        15:32:35 EZB4045I     RESPONSE from 9.67.116.66 -> 520:
                          .
4        15:32:35 EZB4049I         destination 9.67.116.65 metric 1
                          .
4        15:32:35 EZB4049I         destination 10.0.0.0 metric 2
                          .
4        15:32:35 EZB4049I         destination 9.67.112.0 metric 16
1        15:32:35 EZB4029I Tue Jun 28 15:32:35:
1     18 15:32:35 EZB4036I CHANGE metric destination 9.67.112.0, router 9.67.116.66, from 2 to 16
      19 15:32:35 EZB3862I NCP_Delete out to 9.67.116.65:
                           Route to 9.67.112.0, type = Subnet
1        15:32:35 EZB3943I Send dynamic update
1        15:32:35 EZB3950I toall:  requested to skip interface NCSTALU1
1        15:32:35 EZB3951I client 9.67.116.65:  supply 10.68.15.255 -> 0 via TR88
2        15:32:35 EZB4045I     RESPONSE to 10.68.15.255 -> 0:
                          .
4        15:32:35 EZB4049I         destination 9.67.112.0 metric 16
                          .
1        15:32:35 EZB3948I Interface TR92 not up
1        15:32:35 EZB3945I Inhibit dynamic update for 2017537 usec
1        15:32:35 EZB3829I Waiting for incoming packets
                          .
1        15:32:35 EZB3894I Transport from 9.67.116.65:  24 bytes of RIP data
1        15:32:35 EZB4045I     RESPONSE from 10.68.0.88 -> 520:
                          .
4        15:32:35 EZB4049I         destination 9.67.112.0 metric 16
1        15:32:35 EZB3829I Waiting for incoming packets
1        15:32:50 EZB4011I client 9.67.116.65:  30 second timer expired (broadcast)
1        15:32:50 EZB3951I client 9.67.116.65:  supply 9.67.116.66 -> 0 via NCSTALU1
2        15:32:50 EZB4045I     RESPONSE to 9.67.116.66 -> 0:
                          .
4        15:32:50 EZB4049I         destination 9.67.116.66 metric 1
                          .
4        15:32:50 EZB4049I         destination 10.0.0.0 metric 1
                          .
4        15:32:50 EZB4049I         destination 9.67.112.0 metric 16
         15:32:50 EZB3951I client 9.67.116.65:  supply 10.68.15.255 -> 0 via TR88
```

```
2        15:32:50 EZB4045I     RESPONSE to 10.68.15.255 -> 0:
                             .
1        15:32:50 EZB3948I Interface TR92 not up
1        15:32:50 EZB3829I Waiting for incoming packets
:
1        15:36:15 EZB3829I Waiting for incoming packets
1     20 15:36:39 EZB4009I client 9.67.116.65:  5 minute timer expired for route to 9.67.112.0
1        15:36:39 EZB4029I Tue Jun 28 15:36:39:
1     21 15:36:39 EZB4030I DELETE destination 9.67.112.0, router 9.67.116.66, metric 16
1                         flags UP|GATEWAY state SUBNET timer 300
1        15:36:39 EZB4011I client 9.67.116.65:  30 second timer expired (broadcast)
1        15:36:39 EZB3951I client 9.67.116.65:  supply 9.67.116.66 -> 0 via NCSTALU1
2        15:36:39 EZB4045I     RESPONSE to 9.67.116.66 -> 0:
                             .
4        15:36:39 EZB4049I          destination 9.67.116.66 metric 1
                             .
4        15:36:39 EZB4049I          destination 10.0.0.0 metric 1
                             .
1        15:36:39 EZB3951I client 9.67.116.65:  supply 10.68.15.255 -> 0 via TR88
2        15:36:39 EZB4045I     RESPONSE to 10.68.15.255 -> 0:
                             .
4        15:36:39 EZB4049I          destination 9.67.116.66 metric 1
                             .
4        15:36:39 EZB4049I          destination 10.68.0.0 metric 1
:
1     22 15:43:01 EZB3895I Transport from 9.67.116.65:  43 bytes of SNMP data
1     23 15:43:01 EZB4182I SNMP request received from NCP client 9.67.116.65
d        =============== Object data (length=13)
d        0000    2b06 0102 0104 1501
d        0008    0709 4374 4207 39f8
d        0010(16)
d        =============== prefix + address (length=12)
d        0000    2b06 0104 0102 0611
d        0008    0943 7441 4207 39f8
d        0010(16)
d     =============== Inbound SNMP packet (post edit) (length=55)
d     0000 3035 0201 0004 0473
d     0008    6e6d 70a0 2a02 0115
d     0010    0201 0002 0100 301f
d     0018    301d 0619 2b06 0104
d     0020    0102 0611 0943 7441
d     0028    2b06 0102 0104 1501
d     0030    0709 4374 4205 0000
d     0038(56)
d     =============== Sending SNMP request to agent (length=55)
d     0000    3035 0201 0004 0473
d     0008    6e6d 70a0 2a02 0115
d     0010    0201 0002 0100 301f
d     0018    301d 0619 2b06 0104
d     0020    0102 0611 0943 7441
d     0028    2b06 0102 0104 1501
d     0030    0709 4374 4205 00f3
d     0038(56)
1     15:43:01 EZB3829I Waiting for incoming packets
1     15:43:01 EZB4194I SNMP sub-agent received DPI request
d     =============== Received DPI request from SNMP agent (length=69)
d     0000    0043 0201 0101 f14b
d     0008    f34b f64b f14b f44b
d     0010    f14b f24b f64b f1f7
d     0018    4bf9 4bf6 f74b f1f1
d     0020    f64b f6f5 4bf4 f34b
d     0028    f64b f14b f24b f14b
d     0030    f44b f2f1 4bf1 4bf7
d     0038    4bf9 4bf6 f74b f1f1
d     0040    f64b f6f6 0007 2b30
d     0048(72)
```

```
1    15:43:01 EZB4072I SNMP sub-agent:DPI GET request
                        (1.3.6.1.4.1.2.6.17.9.67.116.65.43.6.1.2.1.4.21.1.7.9.67.116.66) received
1    15:43:01 EZB4083I iproutenexthop.9.67.116.66
d    ============== Sending DPI response to SNMP agent (length=77)
d    0000    004b 0201 0105 00f1
d    0008    4bf3 4bf6 4bf1 4bf4
d    0010    4bf1 4bf2 4bf6 4bf1
d    0018    f74b f94b f6f7 4bf1
d    0020    f1f6 4bf6 f54b f4f3
d    0028    4bf6 4bf1 4bf2 4bf1
d    0030    4bf4 4bf2 f14b f14b
d    0038    f74b f94b f6f7 4bf1
d    0040    f1f6 4bf6 f600 8500
d    0048    0409 4374 4149 5f3c
d    0050(80)
1    15:43:01 EZB3829I Waiting for incoming packets
1    15:43:01 EZB4068I SNMP response received from agent 9.67.116.66
d        ============== Received SNMP response from agent (length=59)
d        0000    3039 0201 0004 0473
d        0008    6e6d 70a2 2e02 0115
d        0010    0201 0002 0100 3023
d        0018    3021 0619 2b06 0104
d        0020    0102 0611 0943 7441
d        0028    2b06 0102 0104 1501
d        0030    0709 4374 4240 0409
d        0038    4374 4196 95a2 8540
d        0040(64)
d        ============== Object data (length=25)
d        0000    2b06 0104 0102 0611
d        0008    0943 7441 2b06 0102
d        0010    0104 1501 0709 4374
d        0018    4240 2910 0000 0001
d        0020(32)
d        ============== prefix + address (length=12)
d        0000    2b06 0104 0102 0611
d        0008    0943 7441 2b06 0102
d        0010(16)
d        ============== Outbound SNMP packet (post edit) (length=47)
d        0000    302d 0201 0004 0473
d        0008    6e6d 70a2 2202 0115
d        0010    0201 0002 0100 3017
d        0018    3015 060d 2b06 0102
d        0020    0104 1501 0709 4374
d        0028    4240 0409 4374 4100
d        0030(48)
1        15:43:01 EZB4172I SNMP reply sent to NCP client 9.67.116.66
d        ============== UDP data (length=47)
d        0000    302d 0201 0004 0473
d        0008    6e6d 70a2 2202 0115
d        0010    0201 0002 0100 3017
d        0018    3015 060d 2b06 0102
d        0020    0104 1501 0709 4374
d        0028    4240 0409 4374 4168
d        0030(48)
d        ============== UDP header (length=8)
d        0000    00a1 040e 0037 ec9f
d        0008(8)
d        ============== IP header (length=20)
d        0000    4500 004b 0034 0000
d        0008    0411 a18e 0a44 0058
d        0010    0a44 0001 8002 c12c
d        0018(24)
d        ============== Transport PDU header (length=8)
d        0000    0700 0000 0a44 0058
d        0008(8)
```

```
d        =============== Sending Transport PDU to NCP client (length=84)
d        0000    0700 0000 0a44 0058
d        0008    4500 004b 0034 0000
d        0010    0411 a18e 0a44 0058
d        0018    0a44 0001 00a1 040e
d        0020    0037 ec9f 302d 0201
d        0028    0004 0473 6e6d 70a2
d        0030    2202 0115 0201 0002
d        0038    0100 3017 3015 060d
d        0040    2b06 0102 0104 1501
d        0048    0709 4374 4240 0409
d        0050    4374 4100 0007 3568
d        0058(88)
1        15:43:01 EZB3829I Waiting for incoming packets
:
:
0        15:44:30 EZB3834I ******************************************************
0     24 15:44:30 EZB3890I * Recv:  status from 9.67.116.65
0        15:44:30 EZB3891I * Interface: 10.68.0.88 is now inactive - TR88
0        15:44:30 EZB3834I ******************************************************
3     25 15:44:30 EZB4038I *** Packet history for interface TR88 ***
3        15:44:30 EZB4044I Output: trace:
3        15:44:30 EZB4045I     RESPONSE to 10.68.15.255 -> 0:
                            .
3        15:44:30 EZB4049I         destination 9.67.116.66 metric 1
                            .
3        15:44:30 EZB4049I         destination 10.68.0.0 metric 1
                            .
3        15:44:30 EZB4049I         destination 9.67.112.0 metric 2
3        15:44:30 EZB4045I     RESPONSE to 10.68.15.255 -> 0:
                            .
3        15:44:30 EZB4049I         destination 9.67.116.66 metric 1
                            .
3        15:44:30 EZB4049I         destination 10.68.0.0 metric 1
                            .
3        15:44:30 EZB4049I         destination 9.67.112.0 metric 2
:
:
3        15:44:30 EZB4045I     RESPONSE to 10.68.15.255 -> 0:
                            .
3        15:44:30 EZB4049I         destination 9.67.116.66 metric 1
                            .
3        15:44:30 EZB4049I         destination 10.68.0.0 metric 1
3        15:44:30 EZB4044I Input: trace:
3        15:44:30 EZB4045I     RESPONSE from 10.68.0.88 -> 520:
                            .
3        15:44:30 EZB4049I         destination 9.67.116.66 metric 1
                            .
3        15:44:30 EZB4049I         destination 10.68.0.0 metric 1
                            .
3        15:44:30 EZB4049I         destination 9.67.112.0 metric 2
                     3        15:44:30 EZB4045I     RESPONSE from 10.68.0.88 -> 520:
                                                  .
                     3        15:44:30 EZB4049I         destination 9.67.116.66 metric 1
                                                  .
                     3        15:44:30 EZB4049I         destination 10.68.0.0 metric 1
                                                  .
                     3        15:44:30 EZB4049I         destination 9.67.112.0 metric 2
                     3        15:44:30 EZB4045I     RESPONSE from 10.68.0.88 -> 520:
                                                  .
                     3        15:44:30 EZB4049I         destination 9.67.116.66 metric 1
                                                  .
                     3        15:44:30 EZB4049I         destination 10.68.0.0 metric 1
                                                  .
                     3        15:44:30 EZB4049I         destination 9.67.112.0 metric 2
                     :
                     3        15:44:30 EZB4045I     RESPONSE from 10.68.0.88 -> 520:
                                                  .
```

```
3        15:44:30 EZB4049I         destination 9.67.116.66 metric 1
                           .
3        15:44:30 EZB4049I         destination 10.68.0.0 metric 1
3        15:44:30 EZB4039I *** End packet history ***
         15:44:31 EZB3829I Waiting for incoming packets
                           .
                           .
                           .
1        15:44:41 EZB3948I Interface TR88 not up
1        15:44:41 EZB3948I Interface TR92 not up
1        15:44:41 EZB3829I Waiting for incoming packets
                           .
                           .
                           .
```

The following information explains the numbered items in the trace:

1 The port number and the service name are defined as 580 and ncprout in the *hlq*.ETC.SERVICES data set for this NCPROUTE server.

2 NCPROUTE is processing the NCPROUTE.PROFILE definitions.

3 NCPROUTE is establishing the connection with the SNMP agent defined in NCPROUTE.PROFILE.

4 The NCP client is starting the hand-shaking process with NCPROUTE. NCPROUTE is establishing a session with the NCP client.

5 NCPROUTE received a list of inactive interfaces from the NCP client.

6 NCPROUTE is initializing its interface tables with interface information from the NCP client.

7 NCPROUTE is adding a route to its interface tables.

8 NCPROUTE is processing the NCP client GATEWAYS data set. The trace shows NCPROUTE server options and no additional gateway definitions.

9 NCPROUTE received a transport datagram from the NCP client.

10 The trace shows the contents of the datagram in hexadecimal followed by a division of the datagram into its parts (transport PDU header, IP header, UDP header, and UDP data).

11 The trace shows that the NCP client 9.67.116.65 received the broadcasted routing tables from adjacent router 9.67.116.66.

12 The UDP data in the datagram contains two routing table entries.

13 NCPROUTE is adding a new route to its tables from the information received in the transport datagram.

14 NCPROUTE is issuing a request to the NCP client to add the route to its tables.

15 The NCP client 30-second timer has expired, so NCPROUTE supplies its routing tables to other routers.

16 NCPROUTE is responding to the request by sending its routing tables to the requesting router for the NCP client.

17 This line shows an inactive state for interface TR92.

18 The NCP client 3-minute timer expired. The client was broadcast as a network unreachable route (in the range metric 16—infinite), so NCPROUTE updates its routing tables for the NCP client.

19 NCPROUTE is deleting the NCP client from its tables.

20 The NCP client five-minute timer has expired for the route to 9.67.112.0.

21 NCPROUTE is deleting the route to 9.67.112.0 from its tables for the NCP client.

22 NCPR received a transport datagram from the SNMP client through NCP client 9.67.116.65.

23 NCPROUTE is processing the SNMP request.

24 NCPROUTE has received a status notification from the NCP client. The interface TR88 has become inactive.

25 The packet history for the interface TR88 is included in the trace because the interface has become inactive.

Chapter 33. Diagnosing X.25 NPSI problems

This topic discusses how to diagnose X.25 NPSI problems and includes the following subtopics:

- "Operation" on page 818
- "Configuration requirements" on page 819
- "Sources of diagnostic information" on page 820
- "X.25 trace examples" on page 820
- "Steps for diagnosing logon problems" on page 823
- "Session hangs" on page 824

The X.25 NPSI server uses an X.25 network or point-to-point X.25 line to transfer TCP/IP traffic. The X.25 NPSI server is a VTAM application running as a started task. Either the NPSI Generalized Access to X.25 Transport Extension (GATE) or Dedicated Access to X.25 Transport Extension (DATE) can be used. GATE is recommended because it allows NPSI to handle more details of error recovery and allows an X.25 physical link to be shared with other functions.

Details of the GATE and DATE programming interfaces are in *X.25 NPSI Host Programming*, and further diagnostic information is in *X.25 NPSI Diagnosis, Customization, and Tuning*.

Specifications for carriage of IP traffic on X.25 networks can be found in:

RFC 877
: A Standard for the Transmission of IP Datagrams Over Public Data Networks

X25.DOC
: Old DDN X.25 specifications from BBN (available by anonymous FTP from nic.ddn.mil in directory netinfo)

RFC 1236
: IP to X.121 Address Mapping for DDN

RFC 1356
: Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode

Figure 103 on page 818 shows the X.25 NPSI environment.

© Copyright IBM Corp. 2000, 2015

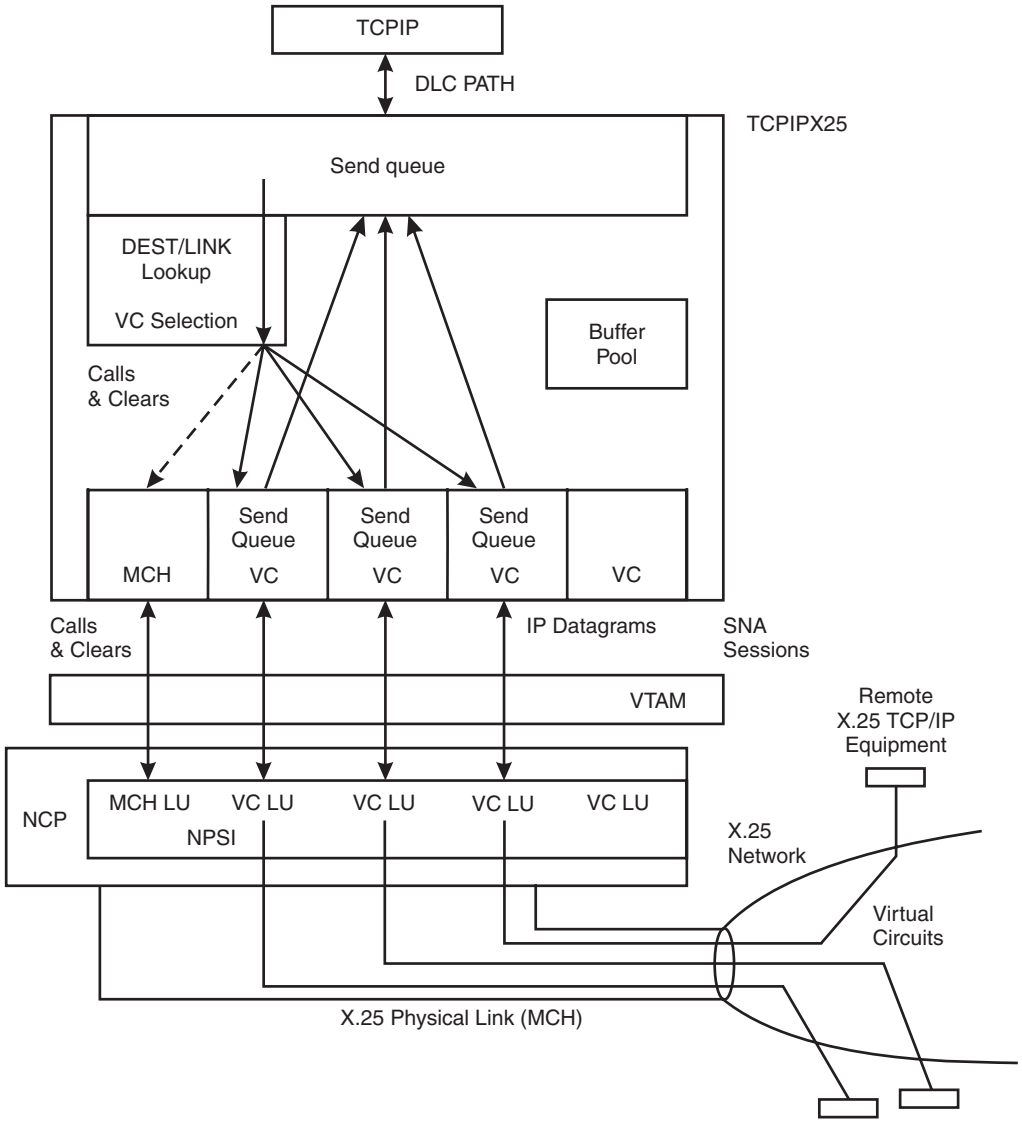

Figure 103. X.25 NPSI environment

Operation

The X.25 NPSI server uses NPSI to set up X.25 virtual circuits as needed to carry traffic to and from remote X.25 equipment. The three main functional areas shown in Figure 103 are:

- TCP/IP interface
- NPSI interface
- IP/X.25 address mapping

IP datagrams are transferred between TCP/IP and the X.25 NPSI server on a DLC path established when a TCPIP X25NPSI device is started. The transfer protocol is similar to that used with SNALINK, with the addition of a first-hop IP address

passed by TCP/IP from the relevant BEGINROUTES statement. The X.25 NPSI server uses the first hop IP address to look up an X.25 address in its destination table.

Communication with NPSI is by way of several SNA sessions. One control session is established at initialization for each MCH LU defined in a LINK statement in the X.25 NPSI server configuration data set. Commands to establish and terminate X.25 virtual circuit connections pass between the X.25 NPSI server and NPSI on the control session. See *X.25 NPSI Host Programming* for details of the control commands. As new virtual circuits are established, NPSI initiates new SNA sessions with the X.25 NPSI server application by means of VTAM LOGON. IP datagrams are then exchanged with the remote equipment over the VC session until an idle timeout occurs or the VC is taken for another destination.

IP addresses are mapped to X.25 addresses by table lookup, or in the case of the DDN network, by a calculation described in RFC 1236. The X.25 NPSI server performs the lookup with the first-hop IP address on each datagram it receives from TCP/IP. The LINK and DEST entries defined in the X.25 NPSI server configuration data set are scanned in order from top to bottom to find a DEST with a matching IP address. After the DEST is found, the link it applies to is selected to carry the datagram, and the active virtual circuits on that link are scanned to find one with an X.25 address that matches the DEST. If such a VC is found, the datagram is queued for transmission on that VC; if none is found and there is a free VC, a new X.25 call is initiated; if all VCs on the link are in use, the least recently used connection is cleared, as long as it has been open for at least the minimum open time, and a new call is initiated. If no VC matches these conditions, the datagram is discarded.

Configuration requirements

The next two topics describe configuration considerations.

RACF/Security Manager requirement

The user ID assigned to the X.25 NPSI start procedure needs an OMVS Segment assigned to it.

VTAM considerations

- APPL definition

  The X.25 NPSI server requires AUTH=(ACQ) and PARSESS=YES in the VTAM APPL definition.
- SWNET definition for switched circuits
  - The value that is specified for MAXDATA for the PU must be at least 10 bytes greater than the value specified for the maximum packet size on the BUFFERS statement in the X.25 NPSI server configuration data set.
  - SSCPFM=USSNTO and DISCNT=(YES,F) are necessary.

NPSI considerations

- BUILD definition

  The value specified for X25.MAXPIU must be at least 10 bytes greater than the value specified for the maximum packet size on the BUFFERS statement in the X.25 NPSI server configuration data set.
- X25.MCH definition

  - LOGAPPL can be coded for recovery.
  - TRAN=NO is required with GATE=DEDICAT.
- X25.VC definition
  - Permanent virtual circuits (PVCs) are not supported.
  - Do not code LOGAPPL except with CONNECT=YES (Fast connect).
  - Do not code MAXDATA except with CONNECT=YES (Fast connect).
- X25.OUFT definition

  X.25 facilities specified with X25.OUFT are not used by the X.25 NPSI server.

Sources of diagnostic information

Many problems with the X.25 NPSI server are the result of configuration faults. Check the following configuration files:

- BEGINROUTES statements, DEVICE and LINK in PROFILE.TCPIP
- The X.25 NPSI server configuration data set
- VTAM APPL definition for the X.25 NPSI server
- NPSI definitions
- VTAM SWNET definitions for NPSI

The primary diagnostic information source is the activity log produced by the X.25 NPSI server. Messages appear in the MVS system log, and can also be captured into a separate data set by including a SYSPRINT DD statement in the X.25 NPSI cataloged procedure. Normal logging records virtual circuit establishment and termination.

Additional information can be recorded about VC activity by setting the TRACE CONTROL option in the X.25 NPSI server configuration data set. This level is sufficient for almost all problem situations; interpretation of the data requires knowledge of X.25 NPSI packet formats. Tracing of the contents of IP datagrams sent to and received from NPSI is provided by the MVS CTRACE option. For details on using the CTRACE option, see Chapter 5, “TCP/IP services traces and IPCS support,” on page 45.

VTAM buffer traces and NPSI X.25 line traces can also be useful in diagnosing difficult problem situations.

You can perform traces on the X.25 LINKNAME using the TCPIP PKTTRACE command or on the SNA LU name using the VTAM **buffer trace** command. See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45, for details about how to use the IP packet trace facility.

X.25 trace examples

The message severity codes (last position of the message ID) are:

I Information (including trace)

W Warning

E Recoverable error

S Recoverable error

T Unrecoverable error

The following example shows normal initialization:

```
EZB2111I VTAM ACB X25IPI1  opened successfully
EZB2210I MCH XU038    packet level ready
EZB2451I IP AS path accepted for job name TCPIPTES
```

Initialization has four main steps:

1. The configuration file is read and processed.
2. VTAM control blocks are initialized (EZB2111I).
3. NPSI physical links (MCHs) configured by LINK statements are initialized (EZB2210I).
4. TCP/IP establishes a DLC path to the X.25 NPSI server (EZB2451I).

Normal incoming call, TRACE OFF

The following example illustrates a normal incoming call with TRACE OFF:

```
EZB2301I VC F001XU038    incoming call from 00000039 user data CC
EZB2325I VC F001XU038    facilities: pkt1024.
EZB2320I VC F001XU038    NPSI logon LU VL038001
EZB2330I VC F001XU038    call complete
    ...some time later...
EZB2350I VC F001XU038    call cleared, cause=00 diagnostic=C5
EZB2351I VC F001XU038    connection terminated for 00000039: sent 1 received
1 dropped 0
EZB2352I VC 010 closed
```

Notes:

1. The VC identifier F001XU038 ties together the events associated with a single virtual circuit. Messages for one VC are intermixed with messages for other VCs.
2. The X.25 address originating the call (00000039) is reported in the EZB2301I message.
3. X.25 calls can optionally request facilities to be applied, such as window size, packet size, throughput class, and reverse charging. These are reported in the EZB2325I message.
4. EZB2330I "call complete" indicates the virtual circuit is ready for transferring TCP/IP data.
5. An X.25 call can be closed by the originator, the acceptor, or the X.25 network. The cause and diagnostic codes in the EZB2350I message indicate the reason. In the example, cause=00 indicates the originator has closed the connection. Lists of cause and diagnostic codes can be found in *X.25 NPSI Diagnosis, Customization, and Tuning*.
6. EZB2351I reports the number of IP datagrams transferred on the virtual circuit.
7. After the EZB2352I "closed" message is issued, the virtual circuit is ready for reuse by another incoming call or to originate a new call.

Normal incoming call, TRACE DATA

The following example illustrates a normal incoming call with TRACE DATA:

```
EZB2230I MCH XU038    packet received (length=17)
EZB2000I 0000 .0.h..... 0BF00188 00000038 00000039 03420A0A
EZB2000I 0010 .               CC
EZB2301I VC F001XU038    incoming call from 00000039 user data CC
EZB2325I VC F001XU038    facilities: pkt1024.
EZB2302I VC F001XU038    call accept packet sent (length=6)
EZB2000I 0000 .0....          0FF00102 0400
EZB2320I VC F001XU038    NPSI logon LU VL038001
EZB2330I VC F001XU038    call complete
```

```
EZB2332I VC F001XU038    data received (length=276)
EZB2000I 0000 E.......<.....}& 45000114 00100000 3C017F82 820FFD26
EZB2000I 0010 ..}...*k.:wxr-(. 820FFD11 0800AA6B 00BAF778 72ADA88E
EZB2000I 0020 0}.f9kq.,.PF;._n 307D0C66 B96BF118 AC085046 3B83DF6E
           ....data omitted for brevity...
EZB2000I 0110 =_3.             BD5F339D
EZB2331I VC F001XU038    data sent (length=277)
EZB2000I 0000 .E.......<.....} 00450001 14001000 003C017F 82820FFD
EZB2000I 0010 ...}&;.2k.:wxr-( 11820FFD 260000B2 6B00BAF7 7872ADA8
EZB2000I 0020 .0}.f9kq.,.PF;._ 8E307D0C 66B96BF1 18AC0850 463B83DF
           ....data omitted for brevity...
EZB2000I 0110 _=_3.            5FBD5F33 9D

EZB2336I VC F001XU038    inactivity timer expired
EZB2353I VC F001XU038    clear request packet sent (length=5)
EZB2000I 0000 .....            00011300 00
EZB2365I VC F001XU038    clear sent
EZB2333I VC F001XU038    packet received (length=1)
EZB2000I 0000 .                17
EZB2358I VC F001XU038    clear confirmed
EZB2351I VC F001XU038    connection terminated for 00000039: sent 1
                         received 1 dropped 0
EZB2352I VC 010 closed
```

TRACE DATA can be used to record the full contents of IP datagrams as they pass through the X.25 NPSI server. The IP header begins at byte 45 (X'2D') within the IP packet. A reduced trace given by TRACE CONTROL shows only the X.25 control packets (call request, call accept, clear request, and clear confirm). See *X.25 NPSI Host Programming* for the detailed packet formats.

Normal outgoing call, TRACE CONTROL

The following example illustrates a normal outgoing call with TRACE CONTROL:

```
EZB2310I VC F810XU038    outgoing call to 00000039
EZB2311I VC F810XU038    call request packet sent (length=20)
EZB2000I 0000 ......h..... 0B081002 04008800 00003900 00003803
EZB2000I 0010 ....             420A0ACC
EZB2230I MCH XU038    packet received (length=5)
EZB2000I 0000 ...0.            0F0810F0 01
EZB2314I VC 0810XU038    call accepted by  user data
EZB2320I VC 0810XU038    NPSI logon LU VL038001
EZB2330I VC 0810XU038    call complete

EZB2336I VC 0810XU038    inactivity timer expired
EZB2353I VC 0810XU038    clear request packet sent (length=5)
EZB2000I 0000 .....            00011300 00
EZB2365I VC 0810XU038    clear sent
EZB2333I VC 0810XU038    packet received (length=1)
EZB2000I 0000 .                17
EZB2358I VC 0810XU038    clear confirmed
EZB2351I VC 0810XU038    connection terminated for 00000039: sent 5
                         received 5 dropped 0
EZB2352I VC 010 closed
```

The steps involved in outgoing and incoming calls are similar. One important difference is that the virtual circuit identifier changes when the call is accepted (compare the EZB2311I and EZB2314I messages). This is related to the details of the NPSI programming interface.

X.25 experts should note that some X.25 packets do not appear in the trace because they are generated by NPSI without the direct involvement of the host application.

Clear confirm is one example. Also, the sequence of events during closing can vary slightly in normal operation, and in some instances, benign VTAM request failures can be reported with message EZB2411E.

Results of LIST command

The following example illustrates the results of the LIST command:

```
EZB2020R MCH XU038    state 1050
EZB2021R VC 010 LU VL038001 DTE 00000039        state 4050
EZB2021R VC 00F LU          DTE                 state 1010
                ...
EZB2021R VC 001 LU          DTE                 state 1010
EZB2022R IP AS TCPIPTES state 80
```

The LIST command is useful to get a snapshot of virtual circuit status. This example shows a normal status with one active VC (state 4050). VC state 1010 indicates ready but not in use. With the NPSI fast connect feature, the normal idle state is 1050. Other intermediate states can appear while an X.25 call or clear is in progress. The codes are listed in z/OS Communications Server: IP Messages Volume 1 (EZA).

The status of the path to TCP/IP is shown in the last line:

- 80 is normal
- 00 indicates that the TCPIP X25 NPSI device has not been started

Termination by TCPIP STOP device

The following example illustrates termination using the TCPIP STOP device:

```
EZB2091I HALT notice accepted, type 0
EZB2250I MCH XU038    terminating
EZB2352I VC 010 closed
EZB2352I VC 00F closed
                ...
EZB2352I VC 001 closed
EZB2480I IP AS TCPIPTES disconnected: sent 7 received 7 dropped 0
EZB2090I Terminating
EZB2099I Ended
```

EZB2480I reports the number of IP datagrams transferred on the DLC path for TCP/IP.

Steps for diagnosing logon problems

Several steps must take place successfully to establish an X.25 virtual circuit for TCP/IP activity:

1. An X.25 call request is received by the X.25 NPSI server from the X.25 network (incoming call) or is sent by the X.25 NPSI server to establish a connection to a new destination (outgoing call).
2. An X.25 call accept confirms the X.25 call request. Call accept is sent by TCPIPX25 for an incoming call, or received from the X.25 network for an outgoing call.
3. NPSI initiates an SNA session with the X.25 NPSI server application by using a VTAM LOGON.

Each of these steps is reported in the activity log, shown in the “X.25 trace examples” on page 820. Problems fall into two main areas: failure of the X.25 call

itself, indicated by either a refusal or an immediate clear, or failure of the NPSI LOGON. Call failures are reported with X.25 cause and diagnostic codes. Standardized cause codes include:

Code	Meaning
00	DTE clearing. The remote system cleared the call.
01	Number busy. The called number cannot accept another call.
03	Invalid facility request. A facility that is requested by the caller is not subscribed or conflicts with a subscribed option.
05	Network congestion. Congestion conditions or some other problem within the network temporarily prevent the requested virtual circuit from being established.
09	Out of order. The called number is out of order.
0B	Access barred. The caller is not permitted to obtain a connection to the called number.
0D	Not obtainable. The called number is not assigned or is no longer assigned.
11	Remote procedure error. An X.25 protocol error at the remote equipment.
13	Local procedure error. An X.25 protocol error.

See *X.25 NPSI Diagnosis, Customization, and Tuning* for a list of diagnostic codes. X.25 networks can also have special diagnostic codes in the range 80–FF.

VC LOGON can fail for various reasons. Among the most common reasons are:

- Incorrect VTAM switched circuit definitions. IDNUM entries are error prone; SSCPFM=USSNTO and DISCNT=(YES,F) are necessary.
- A default VTAM USS table ISTINCDT that is modified to include text in the message 10 entry.
- Coding LOGAPPL on the NPSI X25.VC definitions. Use LOGAPPL only on the X25.MCH and on the X25.VC with the Fast Connect feature.
- Insufficient number of type 1 LUs configured on the NCP LUDRPOOL statement.

A VTAM buffer trace with ID=VTAM helps diagnose the first problem. Collect the following configuration documentation before you contact the IBM Software Support Center. X.25 NPSI server configuration data set, VTAM APPL definition for the NPSI X.25 server, NPSI definitions, and VTAM SWNET definitions for NPSI.

Session hangs

In diagnosing session hang or timeout problems, remember that TCPIPX25 does not track individual TCP sessions; it transfers only IP datagrams. One X.25 virtual circuit can carry datagrams from several TCP sessions. A VC can also be closed and reestablished several times during a TCP session with long periods of inactivity. Failure of an X.25 connection is not directly reflected in TCP sessions it might be carrying, only indirectly by TCP timeouts.

Opening a TCP session, such as a Telnet connection, can fail for reasons not specific to X.25, such as a TCP/IP routing problem caused by an incorrect BEGINROUTES definition, or an IP routing problem in the remote device. Symptoms suggesting these problems include:

- No X.25 call is made when a TCP connection is requested.
- No traffic is received from the remote equipment, indicated by a received count of zero in the EZB2351I connection terminated message.

An established TCP connection can hang because the X.25 network or remote device is down. This is indicated by a clear cause and diagnostic, as described in "Steps for diagnosing logon problems" on page 823.

Helpful hints

PING fails but Telnet and FTP connect. Setting up a new X.25 connection might take longer than the default PING timeout on a busy system. Use the PING TIMEOUT or COUNT parameters to extend the waiting time. Use the NPSI GATE Fast Connect feature to reduce connection setup time.

PING succeeds but Telnet or FTP data transfer times out. Full-screen Telnet and FTP data transfers create large IP datagrams, while PING uses smaller ones. If the small datagrams go through but large ones do not, there might be a problem with MAXDATA on the VTAM switched circuit definitions; see "Configuration requirements" on page 819 for details. Attempting to pass a datagram larger than MAXDATA on a virtual circuit hangs the VC for all subsequent traffic.

A load-dependent hang can be due to an insufficient number of virtual circuits.

The TRAFFIC command can be used to observe virtual circuit data transfer activity.

Documentation requirements

If IBM Support Center help is needed, collect the following configuration documentation before contacting IBM:

- X.25 NSPI server console log showing X.25 connections related to the problem
- X.25 NPSI server configuration data set
- PROFILE.TCPIP data set
- NPSI definitions
- VTAM SWNET definitions for NPSI

Chapter 34. Diagnosing IMS problems

This topic describes how to diagnose IMS problems, and contains the following subtopics:

- "Steps for setting up the IMS TCP/IP services socket interface system" on page 829
- "Common configuration mistakes" on page 830
- "Quick checklist for common problems" on page 831
- "Documentation references for problem diagnosis" on page 844

The IMS TCP/IP Services socket interface allows TCP/IP clients to access IMS using a TCP/IP network. This access is fully described in z/OS Communications Server: IP IMS Sockets Guide. A sockets program-to-program connection is established between a client (TCP/IP socket) program and a server (IMS application) program. TCP/IP and the Listener are agents in the connection establishment. The components of the IMS TCP/IP socket interface system are shown in Figure 104 on page 828.

© Copyright IBM Corp. 2000, 2015

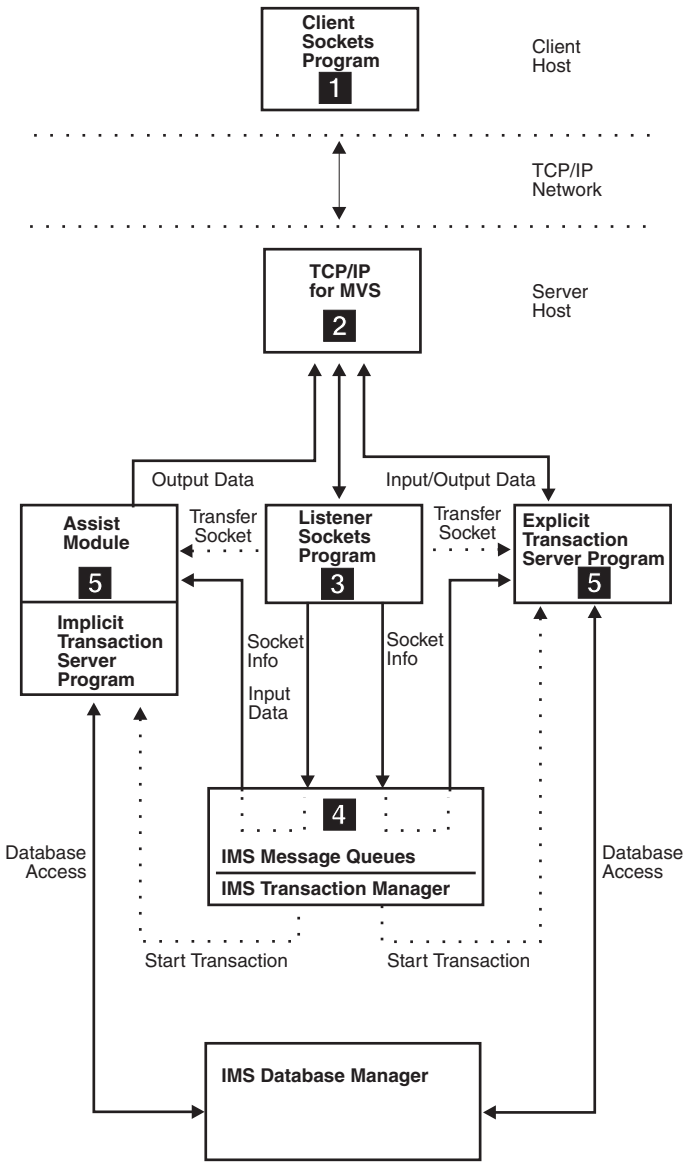

Figure 104. Components of the IMS TCP/IP services socket interface system

The following list is a brief description of the component interaction and data flow that occurs when a client program requests an IMS transaction.

1 The client program starts and sends the transaction request message (TRM) to the Listener port.

2 The Listener reads the TRM and accepts the socket connection between the client program and the Listener from TCP/IP.

3 The Listener validates the TRM, prepares to give the socket connection to the IMS transaction, builds the transaction initiation message (TIM) containing the socket connection information, and sends the TIM to the IMS transaction manager message queue. For implicit IMS transactions, the Listener also reads the input data from the client program and sends it to the message queue.

4 The IMS transaction manager schedules the requested transaction.

5 IMS Transaction. This can be one of the following:

Implicit
The IMS assist module receives the TIM on behalf of the implicit IMS transaction and takes the socket connection from the Listener. The input data is read and the IMS transaction performs the required database access. The IMS assist module, on behalf of the implicit IMS transaction, writes the output data to the client program, through the socket connection, followed by the commit status message (CSM). The socket connection then closes.

Explicit
The explicit IMS transaction receives the TIM and takes the socket connection from the Listener. Input and output data is read and written as defined by the protocol, and the required database access is performed. The explicit IMS transaction writes the CSM to the client program and closes the socket connection.

The IMS transaction and the client program terminate.

Steps for setting up the IMS TCP/IP services socket interface system

This list of steps can be used to diagnose problems in starting components by identifying the prerequisites.

About this task

The steps immediately preceding a step in which you are told to start a component are required to give definitions and configuration information that must be completed correctly before that component can be started. The reference keys in the steps refer to the components as shown in Figure 104 on page 828. All components except the client sockets program belong to the server host.

Procedure

Perform the following steps to establish the system that is described in Figure 104 on page 828.

1. Configure TCP/IP to reserve the Listener port number.

   A TCP/IP port should be reserved for the Listener to connect to when it starts. The following is a sample profile statement to reserve the Listener port.

   ```
   PORT 4096 TCP EZAIMSLN
   ```

   See z/OS Communications Server: IP IMS Sockets Guide for details about the PORT statement.
2. Configure the TCP/IP network from the server host to the client host.

   For the client program to issue IMS transaction requests across a socket connection, there must be a TCP/IP network defined between the client and server hosts. Any physical network supported by IBM MVS TCP/IP can be used to establish this socket connection.

   See the appropriate information in z/OS Communications Server: IP Configuration Reference for details about how to configure the required network to the server host TCP/IP.
3. **2** Start the TCP/IP address space on the server host
4. Establish and verify the network connection from the client host to the server host.

Depending on the network connection, start or activate the required device drivers and network nodes required to establish a TCP/IP network connection.

To verify the TCP/IP network connection, use the PING command on the client host, using the server host destination IP address or network name.

5. Define the Listener to the IMS transaction manager.

   The IMS transaction manager must be defined to expect message queue input from the Listener. For information about how to define the Listener to IMS, see the Listener IMS definitions in z/OS Communications Server: IP IMS Sockets Guide.

6. **5** If the IMS transaction that is requested by the client program is not already written, write it.

   See z/OS Communications Server: IP IMS Sockets Guide for specific details about writing IMS transactions that can be requested by a TCP/IP client program.

7. Define the IMS transaction that is requested by the client program to the IMS transaction manager.

   The IMS transaction must be defined to IMS before the Listener can request it to be scheduled on behalf of the client program. See z/OS Communications Server: IP IMS Sockets Guide for important restrictions when defining IMS transactions.

8. **4** Start the IMS transaction manager and the IMS database manager.

9. Complete the Listener configuration data set.

   The Listener configuration data set is read when the Listener is started. The procedure used to start the Listener (usually EZAIMSLN) uses the ddname LSTNCFG to specify the Listener configuration data set. Following is an example statement that specifies TCPIP.LISTENER.DATA as the configuration data set.

   ```
   LSTNCFG  DD   DSN=TCPIP.LISTENER.DATA,DISP=SHR
   ```

   This data set must contain a minimum set of required statements to specify the environment the Listener is started in and the list of IMS transactions available to client programs.

   See z/OS Communications Server: IP IMS Sockets Guide for details about the format and contents of this data set.

10. **3** Start the Listener address space.

    The Listener is started as an MVS address space as described in z/OS Communications Server: IP IMS Sockets Guide. The JCL procedure required for starting the address space is also listed in z/OS Communications Server: IP IMS Sockets Guide.

11. Write the client program, if not already written.

    See z/OS Communications Server: IP IMS Sockets Guide for programming details about client programs that can request IMS transactions over a TCP/IP network.

12. **1** Start the client program.

Common configuration mistakes

The following is a list of common configuration mistakes:

- The IMS transaction has not been defined in the Listener configuration data set.
- The Implicit or Explicit parameter in the Listener configuration data set does not match the protocol used by the IMS transaction.

- The program specification block (PSB) for the Listener does not include the ALTPCB label.
- The IMS transaction invoked by the Listener does not specify the MODE=SNGL parameter on the IMS TRANSACT macro in the IMS database manager definition. See the z/OS Communications Server: IP IMS Sockets Guide for information about restrictions on application programs.
- The IMS transaction invoked by the Listener was not defined to the IMS transaction manager as a multisegment transaction.
- The IMS transaction invoked by the Listener is an IMS conversational transaction or executes in a remote Multiple Systems Coupling (MSC) environment.

Quick checklist for common problems

The following list summarizes some initial checks that can be made quickly and are helpful in identifying the problem area.

1. Is the TCP/IP network active?

   To verify that the network to the server host is active, use the PING command on the client host, using the same IP address or host name as specified in the client program.
2. Is the Listener started and active on the server host?

   Check that the Listener address space is active and running. The MVS SDSF facility can be used to view the active address space list. Also see "Using NETSTAT" on page 846 for details about how to determine whether the Listener TCP/IP port is active.
3. Did the Listener program list any configuration errors to the SYSPRINT data set?

   Check the JCL DD statement in the Listener start procedure to identify the destination of the SYSPRINT output. See "Where to find error message documentation" on page 847 to determine the reason for any errors. The Listener address space might need to be stopped to flush any error messages to the destination.
4. Have you completed all of the required definitions. See "Steps for setting up the IMS TCP/IP services socket interface system" on page 829 for the list of required definitions and configurations.
5. Is the client program connecting to the same TCP/IP port as the Listener? See "Using NETSTAT" on page 846 for details about how to use the NETSTAT command to identify which port the Listener is connected to and which port the client program is establishing a socket connection on.

Component problems

Table 83 on page 832 lists some of the problems related to starting or stopping one of the components in the IMS TCP/IP Services socket interface system.

Table 83. Component problems

Problem	Cause	Resolution
The Listener terminates on startup	1. Incorrect configuration data set. 2. The prerequisites for starting the Listener have not been completed. 3. Incorrect method of starting.	1. Check for configuration error messages written to the SYSPRINT data set and correct the problems (if any). 2. Complete the required steps listed in "Steps for setting up the IMS TCP/IP services socket interface system" on page 829. 3. Ensure the Listener is being started as an MVS address space as described in z/OS Communications Server: IP IMS Sockets Guide. The JCL procedure required for starting the address space is also listed in z/OS Communications Server: IP IMS Sockets Guide.
The Listener does not terminate	The Listener waits for all of the currently open socket connections to close before it responds to the user termination request. If any of the socket connections have hung, the Listener needs to be forcibly terminated.	Force the Listener to terminate using the command specified in the section about stopping the IMS Listener in z/OS Communications Server: IP IMS Sockets Guide. See "Connection problems" on page 833 for a description of how socket connections can hang.
As the Listener is starting, messages are written to the system console asking if IMS should be started	The IMS system should be started before the Listener. If the Listener is started first, the operator is prompted to start the IMS system.	Reply to the console messages to start IMS.
An implicit IMS transaction written in C is experiencing unexpected problems at startup	If IMS transaction programs written in C are not built correctly, the IMS interface fails on startup.	Build the C program correctly as specified in the section about writing an IMS TCP/IP Services server program in z/OS Communications Server: IP IMS Sockets Guide.

Table 83. Component problems (continued)

Problem	Cause	Resolution
The Listener is abending while accepting the TRM	If a user-defined security exit has been linked into the Listener, it might be causing the problem. The security exit is called when validating the TRM. If the security exit has not been written to accept the required linkage and parameters, the Listener abends because the exit runs in the same address space.	Check that the security exit has been written to accept the linkage and parameters as specified in the section on the IMS security exit in z/OS Communications Server: IP IMS Sockets Guide.

Connection problems

Table 84 lists some problems related to the TCP/IP socket connection. They include problems with establishing the connection, transferring data over the connection, and unexpected loss of the connection.

Table 84. Connection problems

Problem	Cause	Resolution
The client program is experiencing intermittent reject connect responses from TCP/IP	The TCP/IP sockets facility has a connection request backlog queue. While this queue is full, further connection attempts are rejected by TCP/IP. Under load, this queue can temporarily fill, causing some client program requests to be silently ignored.	To reduce the frequency of this problem, increase the size of the backlog queue. The size of the queue is controlled by a parameter in the Listener configuration data set and is limited by the SOMAXCONN statement in the TCPIP PROFILE.
The TCP/IP socket connection to the client program is being broken immediately after the implicit IMS transaction is scheduled	The Listener configuration data set might incorrectly define the implicit IMS transaction as explicit. In this case, the Listener does not pass the input data to the IMS transaction through the message queue as expected. The transaction starts, and upon detecting no data, immediately close the TCP/IP socket connection and terminate.	Verify that the TRANSACTION statements in the Listener configuration data set specify the TYPE parameter correctly.

Table 84. Connection problems (continued)

Problem	Cause	Resolution
Connection lockup for an implicit IMS transaction A connection lockup occurs when both the implicit IMS transaction and the client program are waiting for data from the other end of the socket connection.	The Listener might be waiting for the end-of-message (EOM) segment from the client program. The client program must send a valid EOM segment before the Listener instructs the IMS transaction manager to schedule the IMS transaction. If the client program does not send a recognized EOM segment, the Listener waits indefinitely for it, while the client program waits for a response.	Use the IP packet trace facility to determine whether the client program is sending a valid EOM segment. See "Using IP packet trace" on page 845 for details about the IP packet trace facility. See the information about implicit-mode application data in z/OS Communications Server: IP IMS Sockets Guide for the format of the EOM segment.
Connection lockup for an explicit IMS transaction A connection lockup occurs when both the explicit IMS transaction and the client program are waiting for data from the other end of the socket connection.	1. Because the explicit IMS transaction protocol is user-defined, programming errors can easily lead to connection deadlocks. That is, the server is waiting for more data while the client is waiting for a response, and both wait indefinitely. 2. The Listener configuration data set might incorrectly define the explicit IMS transaction as implicit. In this case the Listener waits for valid implicit data from the client program, or if valid data is received, the explicit IMS transaction waits for data from the client program because the Listener has already read the data and written it to the message queue.	1. Use the IP packet trace facility to identify which part of the protocol is failing. See "Using IP packet trace" on page 845 for details about the IP packet trace facility. 2. Verify that the TRANSACTION statements in the Listener configuration data set specify the TYPE parameter correctly. Timeouts, especially in the client program, are recommended when issuing socket READs to avoid deadlocks and allow easy diagnosis. See the information about SELECT calls in z/OS Communications Server: IP IMS Sockets Guide for more information about specifying timeouts for READs.

Table 84. Connection problems (continued)

Problem	Cause	Resolution
Connection lockup for either an explicit or implicit IMS transaction A connection lockup occurs when both the IMS transaction and the client program are waiting for data from the other end of the socket connection.	1. If the TRM sent by the client program is incomplete, the Listener waits indefinitely for the rest of the message. 2. If the IMS transaction does not successfully issue the takesocket to gain the connection from the Listener, the Listener waits for this event indefinitely. The takesocket might not be issued successfully due to one of the following reasons: • The IMS transaction is defined to run in a message processing region that is not started. In this case, the IMS transaction is never scheduled and, therefore, never issue the takesocket. • One of the several TCP/IP socket calls, up to and including the takesocket, might fail and terminate the IMS transaction. • An IMS error can stop the transaction from being successfully scheduled, or, especially in the explicit case, can cause the IMS transaction to terminate before the takesocket is issued.	1. Check the length and format of the TRM by using the IP packet trace facility as described in "Using IP packet trace" on page 845. 2. Check that the IMS transaction is being successfully scheduled by the IMS transaction manager and ensure that any IMS and socket calls issued by the IMS transaction are checked for unsuccessful return codes.

Table 84. Connection problems (continued)

Problem	Cause	Resolution
The takesocket call issued by the IMS transaction fails **Note:** For implicit transactions, the IMS assist module routines issue a takesocket for the first get unique (GU) issued by the transaction. If the takesocket fails, the GU returns ZZ.	1. IMS can, for recovery reasons, abend a transaction and start it again. If the transaction is abended after it has gained the socket connection (through a takesocket call), the TCP/IP socket connection is lost. Although IMS restores the message queue when it restarts the transaction, the takesocket issued by the transaction fails as the socket connection has already been taken from the Listener. 2. An IMS transaction not defined as multisegment to the IMS transaction manager is scheduled as soon as the TIM is added to the message queue. This gives the IMS transaction an opportunity to issue the takesocket before the givesocket is issued by the Listener. The takesocket fails with an error return code.	1. Restart the client program. To reduce the frequency of this problem, determine why IMS is restarting the IMS transaction by using the IMS trace facility. See "IMS traces" on page 846. 2. Make certain the IMS transaction is defined as multisegment.
The client program is always receiving reject connect responses from TCP/IP	The maximum number of active sockets might have been reached, with all the currently active socket connections unable to complete. An increasing number of socket connections eventually reduces the available socket connections to zero when the number of socket connections equals the MaxActiveSockets configured for the Listener. When this happens, TRMs are not processed by the Listener, and they are left on the TCP/IP backlog queue. When the backlog queue fills, TCP/IP silently ignores a client program connection attempt.	Identify the client programs causing the problem using the NETSTAT command as specified in "Using NETSTAT" on page 846; then continue diagnosis to determine why these connections are locking up. The Listener must be restarted to clear the active socket list. Because there are active socket connections, the Listener must be forced to terminate using the command specified in z/OS Communications Server: IP IMS Sockets Guide.

Table 84. Connection problems (continued)

Problem	Cause	Resolution
Connection lockup or loss when passing a socket connection from one explicit IMS transaction to another A connection lockup is when the socket connection reaches a state where it never completes.	To pass a socket connection from the first IMS transaction to the second, the first IMS transaction must wait after it issues the givesocket until the second IMS transaction issues a takesocket; otherwise, the connection is lost. A connection lockup can occur when the first IMS transaction waits for the takesocket to be issued, but both IMS transactions are defined to run in the same message processing region. In this case, they cannot both be scheduled to run at the same time, and the first IMS transaction waits indefinitely for the takesocket from the second IMS transaction, which is never scheduled.	When passing a socket connection between IMS transactions, make sure the first transaction waits for the second to issue the takesocket and that both IMS transactions can be scheduled to run at the same time.

Error message and return code problems

Table 85 lists problems related to error responses.

Table 85. Error message and return code problems

Problem	Cause	Resolution
The client program is receiving a request status message (RSM)	The Listener sends this message to the client program when it detects an error condition.	Use the return and reason codes from the message to look up the explanation. See "Where to find return code documentation" on page 847.
The implicit IMS transaction is receiving return codes in the I/O program communication block (PCB) that are not defined in the section on status codes in	The IMS assist module performs several socket-related functions on behalf of the implicit IMS transaction in response to IMS transaction manager requests. When errors are detected that are not related to the IMS transaction manager request, the IMS assist module sets special return codes in the PCB.	Look up the meaning of the special return codes. See "Where to find return code documentation" on page 847.

Table 85. Error message and return code problems (continued)

Problem	Cause	Resolution
The Listener error messages are written to the MVS system console instead of the SYSPRINT data set	If the Listener experiences data set I/O errors, it redirects the error messages to the MVS system console.	Check the MVS system console log for I/O errors on the data set to identify the problem. The SYSPRINT DD statement in the JCL procedure to start the Listener specifies the destination data set for the error messages.

Socket data protocol problems

Table 86 lists problems related to data transfer over the socket connection. They include incorrect data sent, not enough or too much data sent, and data corruption.

Table 86. Socket data protocol problems

Problem	Cause	Resolution
The Listener is not responding to the client program	1. If the TRM sent by the client program is incomplete, the Listener waits indefinitely for the rest of the message. 2. If the port specified by the client program is not the port that is attached to the Listener, and the socket connection is established, the other end of the connection does not communicate with the client program as required.	1. Check the length and format of the TRM by using the IP packet trace facility as described in "Using IP packet trace" on page 845. 2. Check that the Listener is attached to the port used by the client program to establish the socket connection. Use the command specified in "Using NETSTAT" on page 846.
All the input data sent from the client program is not being passed to the implicit IMS transaction from the Listener	Any input data written after the first EOM segment is ignored by the Listener.	Check for EOM segments being sent by the client program by using the IP packet trace facility described in "Using IP packet trace" on page 845. See the information about the implicit-mode application data in z/OS Communications Server: IP IMS Sockets Guide for the format of the EOM segment.

Table 86. Socket data protocol problems (continued)

Problem	Cause	Resolution
Explicit IMS transaction is receiving garbled data from or sending garbled data to the client program	The data might need translation when the client program does not exist on an EBCDIC host. For explicit data transfer, the client program, or the IMS transaction, or both, must provide ASCII to EBCDIC translation and byte-order translation of fixed-point binary integers, if required. The Listener automatically translates the TRM when creating the TIM.	Code the client program or the IMS transaction or both to provide the necessary translation when the client program is not on an EBCDIC host.

Table 86. Socket data protocol problems (continued)

Problem	Cause	Resolution
Implicit IMS transaction is receiving garbled data from or sending garbled data to the client program	The automatic data translation when the client program does not exist on an EBCDIC host can be causing the problem. For implicit data transfer, the Listener automatically translates input data from ASCII to EBCDIC, based on the TRM contents. The IMS assist module also automatically translates output data from EBCDIC to ASCII when sending to an ASCII client program, as determined by the TRM. If the TRM sent by the client program is not either ASCII or EBCDIC as required, then the automatic translations fail. The client program is also responsible for any required byte-order translation of fixed-point binary integers. **Notes:** 1. If the data translated between ASCII and EBCDIC contains any nonprintable data, such as integers, flags, or reserved fields, the data is corrupted. In this case, the client program must provide EBCDIC data (including the TRM) for the IMS transaction and expect EBCDIC data from the IMS transaction. 2. If the data is translated between ASCII and EBCDIC and contains characters that are not common to both the ASCII and EBCDIC tables, the nontranslatable characters are translated to spaces.	Code the client program to provide the necessary translation when the client program is not on an EBCDIC host and the automatic data translation cannot be used.
The security exit does not validate user data from the client program	The security exit might not be successfully linked into the Listener. The exit must be compiled and assembled and then linked into the Listener for it to be called.	Check that the security exit has been coded and built correctly as specified in z/OS Communications Server: IP IMS Sockets Guide.

Table 86. Socket data protocol problems (continued)

Problem	Cause	Resolution
Data is corrupted after an implicit IMS transaction issues a GU	The I/O area declared might be too small. When using the IMS assist module, the I/O area provided for the GU call must be large enough to hold the TIM, even though the data eventually returned in the I/O area can be smaller.	Make certain the implicit IMS transaction has enough storage declared to hold the TIM. The size of this message is specified in z/OS Communications Server: IP IMS Sockets Guide.
The PL/I IMS transaction is receiving or sending message segments that are not valid	The message segments might be declared incorrectly. The PL/I API interface to the IMS transaction manager defines the message segments with a four-byte length field, but the length value must include only two of those bytes plus the rest of the segment.	Use the following rules to avoid problems: • The IMS assist module PL/I API routines mimic the interface used by the PL/I API routines. Code PL/I implicit transaction message segments in exactly the same manner as for this interface. • Code the client program in exactly the same manner as for all the IMS transaction API interfaces. The IMS assist module routines automatically converts the message segments from the PL/I API to the standard format. • Explicit transactions do not use the IMS assist module. The message segment format, if required, must match on both the client program and the IMS transaction sides. It is recommended that the standard message segment format be used. See the information about programming considerations for the implicit-mode server and the explicit-mode server in z/OS Communications Server: IP IMS Sockets Guide for more details about the PL/I API issues.

IMS transaction build problems

Table 87 on page 842 lists some problems related to building a component in the IMS TCP/IP Services socket interface system.

Table 87. IMS transaction build problems

Problem	Cause	Resolution
Unresolved external reference errors are causing the linker to fail when linking an IMS transaction	1. The implicit IMS transaction link JCL is not including the IMS assist module and the MVS TCP/IP Services sockets library to resolve external references. 2. The explicit IMS transaction link JCL is not including the MVS TCP/IP Services sockets library to resolve external references.	1. Compare the link JCL to the sample provided in the section about JCL for linking an implicit-mode server in z/OS Communications Server: IP IMS Sockets Guide. 2. Compare the link JCL to the sample provided in the section about JCL for linking an explicit-mode server in z/OS Communications Server: IP IMS Sockets Guide.

IMS database problems

Table 88 lists some problems related to unexpected IMS database actions or failures. They include changes not made or requests for changes that fail.

Table 88. IMS database problems

Problem	Cause	Resolution
The IMS transaction is terminating without performing the required function and without issuing any error messages	The IMS transaction might not be checking for interface errors.	It is the responsibility of the IMS transaction programmer to identify and issue error messages if the IMS database manager, IMS transaction manager, or TCP/IP socket interfaces fail.
The client program is not receiving any data from the implicit IMS transaction, but is receiving a successful CSM	The IMS transaction might be issuing an IMS database rollback (ROLB) call. If the IMS transaction issues a ROLB call, all output accumulated by the IMS assist module is discarded as part of the ROLB function. Depending on how the IMS transaction is coded, it might complete without further output (ISRT calls).	Use caution in issuing ROLB calls in implicit IMS transactions using the IMS assist module. Make certain you understand the details about implicit-mode support for ROLB processing in z/OS Communications Server: IP IMS Sockets Guide.

Table 88. IMS database problems (continued)

Problem	Cause	Resolution
Local IMS transaction manager ISRT/GU/GN calls are failing when issued in IMS transactions	Local calls assume a terminal has requested the IMS transaction. The input and output of data, however, is sent across the socket connection for IMS transactions started by the Listener. The following is a list of specific causes of the problem: • The ISRT call has no terminal associated with the IMS transaction for the output. • There is no data on the message queue for explicit IMS transactions to get with the GU or GN calls. • An implicit IMS transaction receives an unexpected TIM in response to a GU call.	Do not issue local IMS transaction manager calls from transactions started by the Listener. An implicit IMS transaction must use the IMS assist module calls, which accesses either a terminal or socket connection, as required. An explicit IMS transaction must interface directly to the socket connection.
The ISRT call fails for an implicit IMS transaction if a large amount of data is output	The IMS assist module restricts the total output for a single IMS transaction execution to 32KB.	Limit the output for an implicit IMS transaction using the IMS assist module to a total of 32KB.

Table 88. IMS database problems (continued)

Problem	Cause	Resolution
The IMS database manager commits the changes made by an IMS transaction, but the client program receives an error	1. The implicit IMS transaction does not issue a second GU. The IMS database commits the changes either when the IMS transaction ends or when another GU is issued. For implicit IMS transactions, the IMS assist module routines sends the output data and CSM to the client program and closes the socket connection when the second GU is issued. If the implicit IMS transaction does not issue another GU, the changes are committed when the transaction ends, but the client program assumes failure when the CSM is not received. 2. The socket connection might have been broken after the changes were committed but before the CSM was sent. In this case, the client program assumes failure, but the changes have been committed.	1. Implicit IMS transactions that are started by the Listener must issue GU calls to get the next transaction request until the GU call returns with no requests to process. 2. Where possible, the client program should be coded to automatically restart the IMS transaction and handle the condition where the IMS transaction is duplicated. For explicit IMS transactions, a more rigorous protocol can be implemented. **Guideline:** This should be considered as an uncommon case.
The client program does not receive a valid CSM from an implicit IMS transaction	The client program might not have completed the response protocol correctly. The client program must read the response data until it reads an EOM segment. The CSM immediately follows the EOM.	Use the IP packet trace facility to determine whether the IMS transaction is sending a valid EOM segment followed by a valid CSM segment. See "Using IP packet trace" on page 845 for details about the IP packet trace facility. If the correct message segments are being sent, correct the client program to receive the response data. See z/OS Communications Server: IP IMS Sockets Guide for the format of the EOM and CSM segments.

Documentation references for problem diagnosis

This topic contains the information and documentation references required to gather and decode diagnostic information about the IMS TCP/IP Services socket interface system.

The two main tools used for problem diagnosis are the IP packet trace facility and the NETSTAT utility. The use of these tools is explained in following sections and example statements and commands are provided. An explanation of how to interpret the output from each of these tools is also provided.

For TCP/IP or IMS-specific tracing, reference is made to the appropriate diagnosis documentation.

Two cross-reference sections, which list all the types of return codes and error messages that can be issued from the IMS TCP/IP Services socket interface system, are provided at the end of this section. For each type of return code and error message, a reference is made to existing documentation that provides a complete description.

Traces

The following traces can be used to gain information about data flows and actions of the IMS TCP/IP Services socket interface system. The IP packet trace facility is the most helpful trace facility when writing and debugging your own client programs and IMS transactions. The TCP/IP internal traces are mainly used to diagnose problems with the TCP/IP network and socket-specific problems. The IMS traces are mainly used to diagnose IMS-specific problems, such as IMS transaction scheduling and database commit and rollback errors. The IMS assist module trace is used to determine problems with the IMS Assist module. This trace can be enabled by adding a sysdebug dd card to the IMS region procedure where the IMS transaction using the Assist Module is running.

Using IP packet trace

Use IP packet trace to identify the flow of data between the client program and the Listener and IMS transaction servers. TCP packets can be traced on the socket connections established through the Listener-reserved port. If the IP address of the client program is specified, only packets originating from or destined to the client program are traced. Specifying this parameter is recommended to avoid tracing a large number of unrelated TCP packets.

Restriction: When using X.25 devices to provide the network to the client program, the IP packet trace facility must be activated from the individual device address spaces. The previous example only activates tracing in the TCP/IP address space.

See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45 for details about how to use the IP packet trace facility.

The packets that contain data display the data in hexadecimal digits and, in this case, their EBCDIC characters. The numeric fields in the message segments can be verified from the hexadecimal representation, while any alphabetic data can be verified from the EBCDIC display.

TCP/IP internal traces

The TCP/IP internal traces are sent to CTRACE. This is a key trace used to determine the success or failure of the socket calls made by the IMS Listener and the IMS transactions. These traces provide information about the internals of the TCP/IP address space. This information can be used to diagnose problems in establishing the network between the client program and the server host or in establishing the socket connections. See Chapter 5, “TCP/IP services traces and IPCS support,” on page 45, for details about how to use the TCP/IP internal tracing facility.

IMS traces

The IMS traces provide information about the internals of the IMS database system. You can use this information to diagnose IMS transaction scheduling problems, IMS transaction manager message queue problems, and database change problems that cause rollbacks or *commit* errors. For more information about IMS traces, see the IMS knowledge center at http://www-01.ibm.com/support/knowledgecenter/SSEPH2/welcome.

Using NETSTAT

This section details how to use NETSTAT to query TCP/IP port usage and the state of socket connections. This command can be used to verify that the Listener is active and has opened the correct port and to diagnose problems with the socket connection between the client program and the Listener or IMS transaction.

Restriction: The client program must have the socket connection open for NETSTAT to query the connection status.

The NETSTAT SOCKETS command displays which ports are open to which address spaces and displays active socket connections and their status. Following is sample output from this command (the output shown is valid for V2R10 and V1R2):

```
 READY
netstat sockets

 MVS TCP/IP NETSTAT CS V2R10        TCPIP Name: TCPCS          12:34:56
 Sockets interface status:
 Type   Bound to                Connected to          State    Conn
 ====   ========                ============          =====    ====
 Name: INETD1    Subtask: 006DB5B8
 Dgram  0.0.0.0..37             *..*                  UDP      00000058
 Dgram  0.0.0.0..13             *..*                  UDP      00000057
 Dgram  0.0.0.0..19             *..*                  UDP      00000056
 Dgram  0.0.0.0..9              *..*                  UDP      00000055
 Dgram  0.0.0.0..7              *..*                  UDP      00000054
 Stream 0.0.0.0..623            0.0.0.0..0            Listen   0000004B
 Stream 0.0.0.0..514            0.0.0.0..0            Listen   0000004D
 Stream 0.0.0.0..513            0.0.0.0..0            Listen   0000004C
 Stream 0.0.0.0..512            0.0.0.0..0            Listen   0000004E
 Stream 0.0.0.0..37             0.0.0.0..0            Listen   00000053
 Stream 0.0.0.0..7              0.0.0.0..0            Listen   0000004F
 Stream 0.0.0.0..13             0.0.0.0..0            Listen   00000052
 Stream 0.0.0.0..19             0.0.0.0..0            Listen   00000051
 Stream 0.0.0.0..9              0.0.0.0..0            Listen   00000050
 Name: OSNMPD    Subtask: 006DBA70
 Dgram  0.0.0.0..161            *..*                  UDP      00000013
 Stream 0.0.0.0..1027           0.0.0.0..0            Listen   00000014
 Name: TCPCS     Subtask: 00000000
 Stream 127.0.0.1..23           127.0.0.1..1033       Estblsh  00000045
 Stream 9.67.113.27..23         9.37.81.207..1096     ClosWait 00000039
 Name: TCPCS     Subtask: 006C57B0
 Stream 0.0.0.0..23             0.0.0.0..0            Listen   00000012
 Name: TCPCS     Subtask: 006D56F0
 Stream 127.0.0.1..1026          127.0.0.1..1025      Establsh 0000000F
 Name: TCPCS     Subtask: 006D5CF0
 Stream 0.0.0.0..1025           0.0.0.0..0            Listen   0000000C
 Stream 127.0.0.1..1025         127.0.0.1..1026       Establsh 00000010
 Name: USER18    Subtask: 006A3400
 Stream 127.0.0.1..1033         127.0.0.1..23         Establsh 00000044

 READY
```

See z/OS Communications Server: IP User's Guide and Commands for more details about the usage, parameters, and output of NETSTAT.

Where to find return code documentation

The following list refers to the appropriate return code documentation for all return codes that are expected in the IMS TCP/IP Services socket interface system.

- To the client from the Listener (request status message).

  See the information about the request status message (RSM) segment in z/OS Communications Server: IP IMS Sockets Guide for the format of the RSM and a description of the return codes.

  Guideline: The RSM with the "IMS transaction unavailable to be started" return code, is returned when the IMS transaction previously abended or failed and the IMS transaction manager marked it as not able to be scheduled.
- To the client from an IMS transaction (CSM).

  The CSM is received by the client program when the transaction is successful. This message implies a successful return code. If this message is not received, the client program must assume that the IMS transaction did not complete successfully.
- To the implicit IMS transaction from the IMS assist module (I/O program communication block).

  See the information about the I/O PCB implicit-mode server in z/OS Communications Server: IP IMS Sockets Guide for the format of the I/O PCB and return code explanations.
- To an implicit or explicit IMS transaction from TCP/IP.

  See the information about error messages and return codes for IMS sockets calls in z/OS Communications Server: IP IMS Sockets Guide.
- To an implicit or explicit IMS transaction from the IMS transaction manager.

  For information about DL/I status codes, return codes, and reason codes, see the IMS knowledge center at http://www-01.ibm.com/support/knowledgecenter/SSEPH2/welcome.
- To an implicit or explicit IMS transaction from the IMS database manager.

  For information about DL/I status codes, return codes, and reason codes, see the IMS knowledge center at http://www-01.ibm.com/support/knowledgecenter/SSEPH2/welcome.

Where to find error message documentation

The following list refers to the appropriate error message documentation for all error messages expected in the IMS TCP/IP Services socket interface system.

- Error messages from the Listener are written to the SYSPRINT ddname data set. See the information about the IMS Listener error messages in z/OS Communications Server: IP IMS Sockets Guide for descriptions of the error messages in this data set.
- Error messages from TCP/IP are written to the SYSERROR and SYSDEBUG data sets. See z/OS Communications Server: IP IMS Sockets Guide for descriptions of the error messages in these data sets.

Chapter 35. Diagnosing VMCF/TNF/IUCV problems

This topic describes how to diagnose VMCF/IUCV problems and restartable VMCF/TNF problems.

Diagnosing restartable VMCF/TNF problems

This topic describes how to diagnose restartable VMCF/TNF problems and contains the following subtopics:

- "VMCF or TNF fail to initialize"
- "Abends" on page 325
- "Steps for diagnosing no response to commands"
- "VMCF or TNF does not stop" on page 850

You can configure virtual machine communication facility (VMCF) and termination notification facility (TNF) in two different ways: as restartable subsystems or as nonrestartable subsystems. For details about configuration, see z/OS Communications Server: IP Configuration Reference.

If you choose restartable VMCF and TNF, you might encounter the problems described in this topic.

Note: For information about common VMCF and TNF problems, see z/OS Communications Server: IP Configuration Guide.

VMCF or TNF fail to initialize

If VMCF or TNF fail to initialize with an OC4 abend, there is probably an installation problem. Check the PPT entries for errors. Some levels of MVS do not flag PPT syntax errors properly.

Abends 0D5 and 0D6

If, after removing a user, the system crashes with abends 0D5 and 0D6, the application is probably still running and using VMCF. Users should not be removed from VMCF or TNF without first terminating the affected user.

Steps for diagnosing no response to commands

If VMCF and TNF do not respond to commands, one or both of the nonrestartable versions of VMCF or TNF are still active.

Procedure

Perform the following steps to stop and restart the subsystems:

1. Stop all VMCF and TNF users.
2. Stop the subsystems using the commands FORCE ARM VMCF and FORCE ARM TNF.
3. Restart using EZAZSSI.

© Copyright IBM Corp. 2000, 2015

VMCF or TNF does not stop

If you are unable to stop VMCF or TNF, users probably still exist in the VMCF and TNF lists. Use the F VMCF,DISPLAY,NAME=* and the F TNF,DISPLAY,NAME=* commands to identify those users who are still active; then either cancel those users or remove them from the lists, using the F VMCF,REMOVE and the F TNF,REMOVE commands.

Diagnosing VMCF/IUCV problems with the TSO MVPXDISP command

The TSO MVPXDISP command is used as a debugging aid to display the state of the connections from some address spaces to the VMCF address space. In addition, the command is used to obtain information about storage utilization for VMCF and IUCV-related buffers, as well as routines supporting the underlying PC functions. This information can be used by the IBM Software Support Center to analyze the state of the VMCF address space.

The TSO MVPXDISP command is used to display information about a connection for a single user ID or started task to the VMCF address space, or all connections can be displayed. The command is also used to obtain information about the storage utilization. MVPXDISP must be an Authorized Program Facility (APF) command.

If you have a user application that is hung, issue the TSO MVPXDISP command and keep the output for help in diagnosing the problem.

userid Specifies the name of a user ID or started task for which you want the information concerning the connection to the VMCF address space.

ISAQ Specifies that you want information pertaining to storage utilization within the VMCF address space.

The parameters are optional. If no parameter is specified, information about all connections to the VMCF address space as well as the storage utilization data is displayed.

Figure 105 on page 851 shows a sample of the output received from issuing the TSO MVPXDISP command with the *userid* parameter. The messages in this sample are only displayed if the PROFILE MSGID option is in effect for the TSO user ID.

```
mvpxdisp smtp
EZY2053I MVPXDISP: User SMTP     Asid 002C.                   *****
EZY2054I MVPXDISP: Data @ 15B20AD8 Sm=FF Cr0=000008E1 Flags=D4.
EZY2055I MVPXDISP: Client of the VMCF address space.
EZY2055I MVPXDISP: Client of SMSG.
EZY2055I MVPXDISP: Client of VMCF.
EZY2056I MVPXDISP:  IUCV mask=F8F8, Pending Ctl=0000, Appl=0000.
EZY2057I MVPXDISP: VMCF: Buf=00182BA0, Len=00000118, Flgs=00 User=     Key= 80.
EZY2059I MVPXDISP: VMCF: Pending count=0 Flags=00000000.
EZY2058I MVPXDISP: IUCV: Connections=0, Max=255.
EZY2065I MVPXDISP: IUCV: Ctl flags=00000000 Appl flags=00000000.

.  .  .  .
.  .  .  .
```

Figure 105. MVPXDISP sample output using the userid parameter

The output from the MVPXDISP command, when it is issued with the *userid* parameter, contains the following information:

User User ID associated with the address space control block (ASCB) owning the connection to the VMCF address space.

Asid Address space ID (ASID) for the user ID.

Data Address of the control block containing extended information about the user ID.

Sm= Saved system mask of the user's address space.

Cr0= Control register 0 of the user's address space.

Flags Control flags describing the state of the connection. The meaning of the flag bits is as follows:

X'80' SMSG is allowed.

X'40' User ID is a client of VMCF.

X'20' User ID is a client of IUCV.

X'10' User ID is a client of the VMCF address space.

X'08' Reserved.

X'04' User had the TRANSWAP field specified when initially made a client of VMCF.

X'02' Reserved.

X'01' Reserved.

Client of *text string*
Up to 4 lines of text that describe the settings of the bit fields from the Flags variable that concern the client status of the connection. Possible values for *text string* are:

- VMCF address space
- SMSG
- VMCF
- IUCV

IUCV mask=
Enable mask used with IUCV communications.

Pending Ctl=
Control pending interrupt mask used with IUCV communications.

Appl= Application pending interrupt mask used with IUCV communications.

Buf= Address of the VMCF user external interrupt buffer.

Len= Length of the VMCF user external interrupt buffer.

Flgs= Control flags associated with the VMCF connection. The meaning of the flag bits is as follows:

X'80' Specific AUTHORIZE was performed.

X'40' Priority messages are allowed.

X'20' Connection is in a quiesced state.

X'1F' Reserved.

User= If a specific AUTHORIZE was performed, the name of the user ID with whom the restricted connection was established; otherwise, a blank field.

Key= User key at the time the connection was initialized.

Pending count=
Count of pending VMCF requests that have been sent.

Flags= Control flags associated with pending VMCF requests. Only byte 0 contains defined bit fields. All bit positions in bytes 1 through 3 are reserved. The meaning of the defined flag bits is as follows:

X'80' IRB is scheduled or running.

X'40' VMCF interrupt might be pending.

X'3F' Reserved.

Connections=
Count of active IUCV connections.

Max= Maximum number of IUCV connections allowed.

Ctl flags=
Control flags associated with pending IUCV requests on the control path. Only byte 0 contains defined bit fields. All bit positions in bytes 1 though 3 are reserved. The meaning of the defined flag bits is as follows:

X'80' IRB is scheduled or running.

X'40' IUCV interrupt might be pending.

X'3F' Reserved.

Appl flags=
Control flags associated with pending IUCV requests on the application path. Only byte 0 contains defined bit fields. All bit positions in bytes 1 through 3 are reserved. The meaning of the defined flag bits is as follows:

X'80' IRB is scheduled or running.

X'40' IUCV interrupt might be pending.

X'3F' Reserved.

Figure 106 on page 853 shows a sample of the output received from issuing the MVPXDISP command with the ISAQ parameter. The messages in this sample are only displayed if the PROFILE MSGID option is in effect for the TSO user ID.

```
mvpxdisp isaq
EZY2060I MVPISAQ : VMCF_CVT VMCF_Dsa Header at 00B70F18
EZY2061I MVPISAQ :  Subpool 231 1st Getmain count 1
EZY2062I MVPISAQ :              2nd Getmain count 0
EZY2063I MVPISAQ :  Frame size 1024 Max asked 1024(00000400)
EZY2060I MVPISAQ : MVPXINF  XI_Dsa   Header at 15B13000
EZY2061I MVPISAQ :  Subpool   0 1st Getmain count 1
EZY2062I MVPISAQ :              2nd Getmain count 0
EZY2063I MVPISAQ :  Frame size 8192 Max asked 8192(00002000)
```

Figure 106. MVPXDISP sample output using the ISAQ parameter

The output from the MVPXDISP command, when it is issued with the ISAQ parameter contains the following information:

VMCF_Dsa Header at
Address of the anchor block for the dynamic storage areas used by the routines supporting the program call (PC) function.

XI_Dsa Header at
Address of the anchor block for the dynamic storage areas used by the VMCF address space while it services cross-memory calls.

Subpool
Subpool number from which the storage frames are allocated.

1st Getmain count
Count of the number of times a request was made for storage from the pool and none was available. It can be viewed as the maximum number of concurrent requests.

2nd Getmain count
Count of the number of times a storage request was made for an area that exceeded the frame size. This value should never by other than zero, since the frame sizes were chosen based on the maximum storage request size that should be made by the various routines.

Frame size
Number of bytes (decimal) allocated by a GETMAIN request.

Max asked
Largest area in bytes (decimal) that has been obtained from the storage pool to satisfy a request by the routines that exploit the storage pool.

Chapter 36. Diagnosing problems with IP CICS sockets

This information describes how to diagnose IP CICS Sockets problems with the Customer Information Control System (CICS) and contains the following subtopics:

- "Diagnostic data"
- "Initialization problems" on page 856
- "CICS sockets application problems" on page 858
- "CICS sockets control blocks" on page 859
- "CICS trace" on page 860

CICS is an IBM licensed program that enables transactions that are entered at remote terminals to be processed concurrently by user-written application programs.

For more information about CICS and solving problems with CICS, see the following information:

- z/OS Communications Server: IP CICS Sockets Guide
- CICS information at http://www-01.ibm.com/software/htp/cics/library/
- z/OS MVS Diagnosis: Tools and Service Aids

Diagnostic data

To diagnose problems with IP CICS Sockets, some or all of the following data might be required:

- Message logs
  - System log
  - Message log at the transient-data destination specified by the ERRORTD IP CICS Sockets TYPE=CICS configuration option
- CICS external-trace data set (auxtrace)

  Tip: Using the CICS Trace Control Facility transaction, CETR, ensure the following CICS trace flags are set to obtain the CICS auxiliary trace:
  - Set the CICS Master User Trace Flag to the value of ON to generate IP CICS Sockets CICS trace records
  - Set the Master System Trace Flag to the value of ON to generate CICS trace records
  - Set the AP component trace level to the value of 1

  Rule: Ensure that CICS tracing is enabled for the IP CICS Socket Interface. If the IP CICS Sockets TYPE=CICS TRACE configuration option is NO then no IP CICS Sockets CICS tracing occurs. Either change the configuration option to enable IP CICS Sockets CICS tracing and then stop and restart the IP CICS Socket Interface or dynamically enable the CICS trace by using the EZAO,START,TRACE command or with the EZAO,SET,CICS transaction specifying TRACE=YES.
- Component trace
  - Engine
  - Physical file system (PFS)

© Copyright IBM Corp. 2000, 2015

  - Socket
  - Socket (SOCKAPI)
  - Transmission control protocol (TCP)
- Dumps
  - CICS address dump, if captured.

    Guideline: Ensure the following CICS environment before recreating a problem and taking a dump:
    - The CICS internal trace is started
    - The Master System trace flag and Master User trace flag is on
    - Standard trace level 1-2 set for the AP component
    - IP CICS Sockets CICS tracing is enabled
  - Supervisor Call (SVC) dump. SVC dumps are also known as *console dumps* or *system dumps*.

    Guideline: For hangs and loops, request an SVC dump of CICS and TCP/IP.
- NETSTAT SOCKET output
- NETSTAT CONN output

Initialization problems

This topic describes some problems you might encounter when attempting to initialize CICS configured to use IP CICS Sockets.

Steps for diagnosing CICS socket interface not initialized

Determine why the CICS socket interface did not initialize.

Procedure

Follow these steps:

1. Issue the EZAO,START,CICS command, and then check that the interface initializes.
   a. If the interface initializes, check that EZACIC20 is in the Program Load Table (DFHPLT).

      Putting EZACIC20 into the PLT allows the CICS Socket Interface to initialize on CICS address startup. See z/OS Communications Server: IP CICS Sockets Guide for more information.
   b. If EZACIC20 is defined in the DFHPLT, check the message logs for failures.
   c. If there are no messages, then start CICS with an auxiliary trace active, IP CICS Sockets CICS tracing enabled, and then request an SVC dump of CICS.
   d. Call the Support Center.
2. Verify that the socket Resource Definition Online (RDO) definitions have been properly installed and that the correct data sets are in the STEPLIB and DFHRPL concatenations.

Steps for diagnosing CICS listener not initialized

Determine why the CICS Listener did not initialize.

Procedure

Perform the following steps:

1. Use the EZAC transaction to verify that the listener is defined in the configuration file.
2. In the configuration-file record for that listener, verify that IMMEDIATE is set to YES, and then verify that the correct APPLID and port number are specified.
3. Verify that the listener is properly defined in a CICS RDO group and that the RDO group is in the proper group list.
4. Check the message logs for failures.
   a. If there are no messages, start CICS with auxtrace active IP CICS Sockets CICS tracing enabled, and then request an SVC dump of CICS.
   b. If there are messages, call the Support Center.
5. If an EZY1292E message was issued, investigate why the CICS socket interface did not initialize. See "Steps for diagnosing CICS socket interface not initialized" on page 856.
6. If an EZY1369E message was issued, investigate why the TCP/IP stack as specified on the IP CICS Sockets interface TCPADDR configuration option did not initialize. See "Steps for diagnosing problems" on page 3 for steps on diagnosing TCP/IP problems.

No CICS sockets messages issued

If no CICS sockets messages (error or informational) were issued, verify that the correct CICS transient-data queue is specified in the EZACICD TYPE=CICS ERRORTD field in the configuration record for the CICS region. A *region* is the CICS address space.

Steps for diagnosing TCP/IP clients unable to connect

Determine why TCP/IP clients are unable to connect.

Procedure

Perform the following steps:

1. Verify that the listener is active by logging on to CICS, and then issue a CEMT I TASK command. Make sure that the listener name appears in the task list.
2. Verify that the listener is listening on the correct port number by issuing a NETSTAT CONN command, and then check that the listener has the correct port in listen status. Verify that clients are trying to connect to this port and to the correct IP address.
3. Check the ERRORTD log and verify that the EZY1291I message has been issued. If it has not been issued, look for messages that indicate a failure.
4. If message EZY1365E is issued, then ensure that the value specified for the MAXFILEPROC is larger than the listener's NUMSOCK value. Also, ensure that the client's user ID FILEPROCMAX setting is appropriately specified. For more information about how MAXFILEPROC affects tuning applications, see z/OS UNIX System Services Planning.

   For more information about the FILEPROCMAX specification, see the documentation provided for the SAF product in use on your system. If using RACF, this can be found in z/OS Security Server RACF Security Administrator's Guide.

Steps for diagnosing child-server transactions not starting

Child-server transactions are transactions that are started by the listener.

Procedure

If child-server transactions are not starting, perform the following steps:

1. Issue a CEMT I TRANSACTION command to verify that the transaction is installed. If it is not installed, a `NOT FND` message is displayed.
2. Issue a CEMT I PROGRAM command to verify that the child-server program is installed.
3. If the transaction or program is not installed, define it in the proper RDO group.
4. Check the message logs for failures.

CICS sockets application problems

This topic describes some of the problems you might encounter with CICS sockets applications.

Steps for diagnosing hung CICS tasks

Determine why CICS application tasks hang.

Procedure

Perform the following steps:

1. While a task is hung, request an SVC dump of CICS and TCP/IP.
2. If the problem can be re-created, re-create with CICS auxtrace and component trace turned on.
3. Issue a NETSTAT SOCKET command to determine whether the task is waiting for a particular socket call to be posted. If it is waiting, you can issue the NETSTAT DROP command to end it.
4. If the application is hung while waiting for the completion of a READ command, consider issuing a SELECT or SELECTEX command prior to the READ command. The SELECT command returns either the number of sockets that are ready to be read or 0 if it times out. The SELECTEX command also returns either the number of sockets that are ready to be read or 0 if it times out and it also returns an ECB or a list of ECBs.

Hung CICS region

If a CICS sockets application program using the Call Instruction API (EZASOKET) is erroneously link-edited without the EZACICAL stub, the entire CICS region might hang while waiting for socket calls to complete. Ensure that EZACICAL is explicitly link-edited with the application.

An EZASOKET call should generate a static call to the EZASOKET entry point within the EZACICAL stub. If the application is not compiled and link edited correctly, the EZASOKET call generates a dynamic call to program EZASOKET, which calls the socket API directly.

Errors on socket calls

If you receive errors on socket calls, note the ERRNO that is received, and then look it up in the portion of z/OS Communications Server: IP CICS Sockets Guide that describes return codes.

A SOCKAPI CTRACE can also help diagnose problems with EZASOKET calls.

CICS shutdown hangs

If an EZY1342I message has been issued, there is a CICS task that has at least one socket open and that is not terminating. You can fix this problem by executing an immediate termination of the CICS socket interface rather than a deferred termination. To execute an immediate termination, issue an EZAO,STOP,CICS command, and then specify YES at the `IMMEDIATE` prompt.

If you do not add EZACIC20 to the shutdown DFHPLT, CICS cannot terminate because the socket subtasks are still attached to the CICS region. To terminate CICS without EZACIC20, manually shut down the CICS socket interface using the EZAO transaction.

If you have added EZACIC20 to the shutdown DFHPLT then set the IP CICS socket interface PLTSDI configuration option to the value YES to force an immediate shutdown.

CICS sockets control blocks

This section describes some problems you might encounter with the task interface element (TIE) and global work area (GWA). For information about the layout of GWA, TIE, and other control blocks, see the section in the z/OS Communications Server: IP CICS Sockets Guide that describes external data structures.

Task interface element

A Task interface element (TIE) represents a CICS task that issued at least one call to the CICS sockets API. You can locate TIEs in a dump of the CICS region by issuing the IPCS VERBX CICSxxx 'UEH=3' command. CICSxxx is the name of the VERBEXIT used to format a CICS TS dump and is specific to the release of CICS TS that produced the dump. After the CICSxxx VERBEXIT returns, then search for EZACIC01.TIE. The CICSxxx EZACIC01 prefix identifies it as a TIE for CICS sockets. For more information about the CICS TS VERBEXITs, see the CICS information at http://www-01.ibm.com/software/htp/cics/library/.

The IPCS VERBX CICSxxx 'UEH=3' command output shows a CICS image of the TIE. The TCP/IP TIE is embedded within the CICS image of the TIE and starts at offset +X'80'.

The IPCS VERBX CICSxxx'UEH=3' command output contains TIEs for other interfaces as well.

Global work area

The GWA is the main anchor point for the CICS socket interface. It contains general status data, work areas, and pointers to other control-block chains. You can locate the GWA in a dump of the CICS region by issuing the IPCS VERBX CICSxxx 'UEH=3' command, and searching for EZACIC01.GWA. The EZACIC01 prefix identifies it as the GWA for CICS sockets.

CICS trace

The CICS sockets task-related user exit (TRUE), EZACIC01, issues CICS trace entries at the following four points of execution:

- When the TRUE receives a socket call from an application
- When the TRUE is passing the socket call to the subtask
- When the TRUE receives the response from the subtask
- When the TRUE is ready to return its response to the application

The trace point ID is AP 00C7. Trace records are self-explanatory. They show the type of call, the point of execution, the ERRNO, and the RETCODE.

Steps for displaying the internal trace

Trace records can be written either to a CICS internal trace table or to its external-trace data set (auxtrace).

Procedure

Perform the following steps to display the internal trace:

1. Request a dump of the CICS region by using the RGN SDATA=(option 1,option 2...option n) parameter on a DUMP command. Examples of options are CSA, PSA, NVC, RGN, TRT, SQA, LSQA, and LPA. For a complete list of options, see z/OS MVS Diagnosis: Tools and Service Aids.
2. Display the trace by using the IPCS VERBX CICSxxx 'UEH=3' command.

   Tip: CICS trace can also be directed to the GTF trace data set.

What to do next

For more information about the auxtrace, see the CICS information at http://www-01.ibm.com/software/htp/cics/library/.

Chapter 37. Diagnosing problems with SMC-R

SMC-R problems are often related to switch configuration, physical network ID (PNetID) configuration, and other configuration issues. Common problems with using SMC-R communication include the following categories:

- "Switch configuration issues"
- "Physical network ID configuration issues" on page 862
- "No associated subnet mask" on page 862
- "PFID status remains STARTING" on page 863
- "Problem with SMC-R interaction with security function" on page 863

The SMCReason field of the Netstat ALL/-A report and the SMCR field of the Netstat DEvlinks/-d report provide information that is related to SMC-R problems. For a complete list of SMCReason codes in the Netstat ALL/-A report and the SMCR Disabled reasons in the Netstat DEvlinks/-d report, see z/OS Communications Server: IP System Administrator's Commands.

Switch configuration issues

Remote Direct Memory Access (RDMA) processing requires standard 10GbE switch support, and distance limitations might exist. Enable the global pause frame (a standard Ethernet switch feature for Ethernet flow control that is described in the IEEE 802.3x standard) on the switch.

When the SMCReason field of the Netstat ALL/-A report is 00005013 - RDMA CONNECTIVITY FAILURE, VTAM was not able to complete the SMC-R Link Confirm flow, which usually indicates a switch configuration issue. The Link Confirm message is the first data sent over the RDMA over Converged Ethernet (RoCE) fabric. Check for the following issues:

- If you are using VLANs, verify that the VLAN configuration on the RoCE switch ports is consistent with the VLAN configuration on the OSD switch ports.

  For example, the OSD switch ports might be configured properly, with no VLAN ID or the default VLAN ID, but the RoCE switch ports have a different VLAN ID configured, such as trunk mode with VLAN IDs 400 and 500.

  For more information about configuring VLANs, see VLANID considerations in z/OS Communications Server: IP Configuration Guide.
- Verify that your cable is plugged into the correct port on the IBM 10GbE RoCE Express feature and into the correct port on the switch.

  For example, the cable might be plugged into the correct port on the 10GbE RoCE Express feature but into the wrong port on the switch, or the cable might be plugged into the correct port on the switch but into the wrong port on the 10GbE RoCE Express feature.
- Verify that the MTU value configured on the switch is large enough to support your configured MTU size on GLOBALTCPIP SMCR MTU for this interface.

  Enable jumbo frame support on the RoCE switch ports.
- Multiple switches are in use but the switch uplinks are not configured properly.

© Copyright IBM Corp. 2000, 2015

- For some switches (for example the IBM RackSwitch™ G8264), you might need to configure the RoCE switch ports as edge ports. This places the port in the forwarding state as soon as the link is up, reducing delays due to Spanning Tree Protocol processing.

Verify that you have Ethernet flow control enabled on your switch. Ethernet flow control is implemented by using pause frames. If the control is not enabled, this can cause the switch to be overrun leading to packet loss.

Physical network ID configuration issues

The TCP/IP stack must be able to determine which physical network is connected to a particular 10GbE RoCE Express interface, so that the 10GbE RoCE Express interface can be associated with the SMC-R capable IPAQENET or IPAQENET6 interfaces that connect to that same physical network.

Use the Netstat DEvlinks/-d and D NET,TRL,TRLE=xxxx commands to verify the physical network ID (PNetID) value on the OSD interfaces and the 10GbE RoCE Express interfaces.

- If the Netstat DEvlinks/-d report for your OSD interface indicates SMCR: DISABLED (NO PNETID), ensure that you configured the PNetID value on the correct OSD port in the HCD definitions.
- If you receive message EZD2028I with reason PNETID IS NOT CONFIGURED during 10GbE RoCE Express interface activation, ensure that you configured the PNetID value on the correct 10GbE RoCE Express port in the HCD definitions.
- If the Netstat DEvlinks/-d report for your OSD interface indicates SMCR: Yes and your 10GbE RoCE Express interfaces initialized successfully, verify that the PNetID value of the OSD interface matches that of the intended 10GbE RoCE Express interfaces.

In the HCD definitions, value PNetID 1 is for port 1 on 10GbE RoCE Express features and port 0 on OSD adapters, and value PNetID 2 is for port 2 on 10GbE RoCE Express features and port 1 on OSD adapters. Values PNetID 3 and PNetID 4 are not used.

For more information about configuring PNetIDs, see Physical network considerations in z/OS Communications Server: IP Configuration Guide.

No associated subnet mask

SMC-R is used only between peers whose IPv4 interfaces have the same subnet value or whose IPv6 interfaces have at least one prefix in common.

- For IPv4, if the SMCR field of the Netstat DEvlinks/-d report for an OSD interface is DISABLED (NO SUBNET MASK), it means no subnet mask value is configured for the OSD interface.
- For IPv4, if the SMCReason code in the Netstat ALL/-A report is 521E PEER SUBNET/PREFIX MISMATCH, the interfaces on the peer stacks use different subnets.
- For IPv6, if the SMCReason code in the Netstat ALL/-A report is 521E PEER SUBNET/PREFIX MISMATCH, the interfaces on the peer stacks do not have a prefix in common.

For information about associating your interfaces with the appropriate subnet or prefix, see Configuring Shared Memory Communications – RDMA in z/OS Communications Server: IP Configuration Guide.

PFID status remains STARTING

The PFIDSTATUS field is the RNIC interface Peripheral Component Interconnect Express® (PCIe) function ID (PFID) status. The following list describes the possible status values:

- READY

  READY indicates that the initialization sequence with the PFID is complete and the PFID is ready.
- NOT ACTIVE

  NOT ACTIVE indicates that the PFID was never started or was stopped after it was started.
- STARTING

  STARTING indicates that a START of the PFID was issued and TCP/IP sent an activation request to the Data Link Control (DLC) layer. This means z/OS Communications Server did not receive a port state change event that indicates the port is active from the 10GbE RoCE Express feature. Until the port state change event is received, the PFIDSTATUS remains in STARTING state.

  Take the following actions if the PFIDSTATUS field does not change from STARTING to READY:
  - Check that your cables are connected properly.
  - Verify that the switch ports are enabled.
  - If the RoCE adapters are hard-wired to each other, the STARTING status is expected until the partner side has started the RNIC interface.
  - Verify that the optical cable used for the RoCE adapter is not damaged.

Problem with SMC-R interaction with security function

Generally, security functions that require TCP/IP to examine TCP packets cannot be used with SMC-R communications because data that is sent over SMC-R links is not converted into TCP packets. For more information, see Security functions in z/OS Communications Server: IP Configuration Guide.

Chapter 38. Diagnosing problems with Express Logon

The Express Logon feature in Communications Server for z/OS allows a user on a workstation, with a TN3270E client and an X.509 certificate, to log on to an SNA application without entering an ID or password.

This topic describes how to diagnose problems using Express Logon for the z/OS Communications Server Express Logon feature, including the Digital Certificate Access Server (DCAS). It contains the following sections:

- "Analyzing start problems with the DCAS" on page 866
- "Analyzing client interface problems" on page 867

For complete information about Express Logon, see the following:

- z/OS Communications Server: IP Configuration Guide
- z/OS Security Server RACF Security Administrator's Guide

For most situations in which the DCAS does not start, a message to the console is displayed. If the explanation in z/OS Communications Server: IP and SNA Codes does not help, you should turn on debugging and logging. You can specify debugging and logging as startup parameters from the z/OS UNIX shell or from the MVS console as a started procedure:

- If the DCAS is started from the z/OS UNIX shell, you can specify the following:

  ```
  dcas -d <debugging_level>  -l <logtype>
  ```

- If the DCAS is started from the MVS console, you can specify debugging and logging on the PARM statement after the final slash, as shown in the following example:

  ```
  //DCAS   PROC
  //*
  //DCAS   EXEC PGM=EZADCDMN,REGION=4096K,TIME=NOLIMIT,
  // PARM='POSIX(ON) ALL31(ON)/-d -1 SYSLOGD'
  ```

The following optional parameters can be used with both DCAS UNIX commands and MVS started procedures:

-d or -D

Indicates debugging. The following levels apply:

1 Specifies log error and warning messages.

2 Specifies log error, warning, and informational messages.

3 Specifies log error, warning, informational, and debug messages.

The default level is 3.

-l or -L

Indicates logging to SYSLOGD or to a designated log file. If you do not specify this parameter, logging defaults to `/tmp/dcas.log`.

If you specify a debug level, but not logging, the DCAS attempts to open the default log file `/tmp/dcas.log`. If this fails, debugging is turned off.

For SYSLOGD, the DCAS uses the log facility local0.

© Copyright IBM Corp. 2000, 2015

If DCAS has already been started you can issue a MODIFY DCAS,DEBUG=*debug_level* from the MVS console to enable, disable, or switch the level of debugging. See z/OS Communications Server: IP System Administrator's Commands for more information about this command.

An accent mark (`) is used in the definition above, not a single quotation mark.

For further aid in diagnosing errors, see the error logs of the TN3270E middle-tier servers. Also, examine the HOD client security message panel.

The following **netstat** commands, issued from the middle-tier server, are useful in determining connectivity problems between z/OS Communications Server and DCAS.

For AIX, the **netstat** command is:

```
netstat -an | grep port#
```

For CS/2, the **netstat** command is:

```
netstat -sn | grep port#
```

For NT, the **netstat** command is:

```
netstat -an | more port#
```

In the **netstat** commands, port# is the listening port of DCAS. The default DCAS port is 8990.

Analyzing start problems with the DCAS

When analyzing problems that occur when starting the DCAS, consider the following:

- The DCAS must run from an APF Authorized library.
- The DCAS uses z/OS Language Environment C runtime services. Make sure that the Language Environment C runtime library is compatible with the current level of z/OS Communications Server.
- If DCAS uses IBM System SSL (TLSMECHANISM DCAS):
  - The DCAS calls SSL cryptographic services runtime library. Verify that *hlq*.SYS1.SIEALNKE is accessible at run time. If certificates are authenticated using the X.500 host, SSL uses LDAP services to access the X.500 host. If running from the z/OS UNIX shell, verify that the LIBPATH environment variable includes /usr/lib.
  - The DCAS attempts to initialize SSL services. If you are using key rings that reside in the z/OS UNIX file system, verify that the KEYRING and STASHFILE keywords in the DCAS configuration file point to valid z/OS UNIX file system file names. Names are case-sensitive. If using key rings that reside in RACF, verify that the SAFKEYRING keyword in the DCAS configuration file references a valid RACF key ring.
- If DCAS uses AT-TLS policies (TLSMECHANISM ATTLS), see the following topics for more information:
  - Customizing DCAS for TLS/SSL in z/OS Communications Server: IP Configuration Guide
  - Chapter 28, “Diagnosing Application Transparent Transport Layer Security (AT-TLS),” on page 707

- The DCAS must be associated with a valid user ID using z/OS UNIX services. It must run with the POSIX(ON) C runtime option. Use the following RACF command:

  ```
  ADDUSER  dcasid   DFLTGRP(OMVSGRP) OMVS(UID(0) HOME('/')
  ```

- If the DCAS is started as an MVS started procedure, verify that the following RACF commands have been issued:

  ```
  RDEFINE STARTED  DCAS.* STDATA(USER(dcasid))
  RDEFINE OPERCMDS (MVS.SERVMGR.DCAS) UACC(NONE)
  PERMIT MVS.SERVMGR.DCAS  CLASS(OPERCMDS) ACCESS(CONTROL) ID(dcasid)
  SETROPTS RACLIST(OPERCMDS)   REFRESH
  ```

- The DCAS uses the TCP/IP protocol to communicate with clients in the network. Verify that the z/OS Communications Server products VTAM and TCP/IP have been started and are active.

Analyzing client interface problems

When analyzing problems with client interfaces, consider the following:

- DCAS uses the TCP/IP protocol to communicate with its clients, the TN3270 middle-tier servers. Verify that the z/OS Communications Server products VTAM and TCP/IP have been started and are active. To verify network connectivity to a client, try pinging that client.
- The DCAS uses RACF services to obtain a user ID given a digital certificate.
  - Verify the certificate has been defined properly to RACF. Use the following commands:

    ```
    SETROPTS CLASSACT(DIGTCERT)
    SETROPTS RACLIST(DIGTCERT)    REFRESH
    PERMIT  IRR.DIGTCERT.function  CLASS(FACILITY)  ID(dcasid)  ACCESS(CONTROL)
    RACDCERT ID(userid)   ADD('certificate dataset name')  TRUST
    ```

  - Verify that the user ID associated with the DCAS has permission to access certificates. Use the following RACF commands:

    ```
    SETOPTS CLASSACT(DIGTCERT)
    SETROPTS RACLIST(DIGTCERT)    REFRESH
    PERMIT  IRR.DIGTCERT.LIST  CLASS(FACILITY)   ID(dcasid)  ACCESS(CONTROL)
    ```

  - The DCAS uses RACF services to obtain a PassTicket for an associated application ID. Verify that the RACF PTKTDATA profile for the application ID has been defined properly. The ID must match the ID specified on the workstation client. For HOD V5, this is the name specified in the Express Logon Application ID pop-up window. It might not be the same name specified on the USSMSG10. For applications such as TSO, specifying the application ID can be difficult since the profile name has special RACF considerations. See the z/OS Security Server RACF Security Administrator's Guide.

    Use these commands to verify the RACF PTKTDATA profile:

    ```
    SETROPTS CLASSACT(PTKTDATA)
    RDEFINE profile PTKTDATA SSIGNON()
    SETROPTS RACLIST(PTKTDATA) REFRESH
    ```

Chapter 39. Diagnosing resolver problems

This topic describes how to diagnose resolver problems and contains the following subtopics:

- "Steps for resolving the hostname"
- "Steps for resolving caching problems" on page 872
- "Steps for responding to message EZZ9308E" on page 875
- "Steps for responding to message EZZ9311E" on page 879
- "TRACE RESOLVER" on page 884
- "CTRACE - RESOLVER" on page 907

The resolver provides two kinds of tracing plus an IPCS subcommand to help analyze resolver problems in dumps. The resolver provides TRACE RESOLVER information that can be helpful in debugging problems an application program could have with using resolver facilities (for example, GetAddrInfo or GetNameInfo). Component Trace is used for tracing the RESOLVER component (SYSTCPRE) for diagnosing resolver problems that cannot be isolated to one particular application. Use the IPCS RESOLVER subcommand to format and summarize resolver control blocks (see "RESOLVER command" on page 282).

See z/OS Communications Server: IP Configuration Reference for additional information.

Steps for resolving the hostname

Determine why a hostname failed to resolve.

Before you begin

You need to know the exact hostname that failed to resolve and the environment in which the application was running (for example, TSO, UNIX, or batch).

Procedure

1. Diagnose why the hostname failed to resolve by pinging the hostname. Base your next course of action on the following conditions:

© Copyright IBM Corp. 2000, 2015

If ping for the hostname...	Then...	Solution
Succeeds, but another application fails when resolving the same hostname	One or more of the following may be the problem: • The resolver configuration for the application in the users environment. • The resolver cache has information about the hostname, and one application is able to access the resolver cache and the other is not. • The resolver cache has saved different information about the hostname as provided by different name servers, and the applications are using different cached information.	Use the Trace Resolver to solve the problem.
Fails, but the hostname is converted to an IP address	The resolution is successful but the host is not reachable or active.	See Chapter 4, "Diagnosing network connectivity problems," on page 25 to continue researching the problem.
Fails to convert the name to an IP address	The problem might be with the resolver configuration, querying the resolver cache, searching local host files, or using DNS.	Use Trace Resolver to solve the problem. **Note:** You can use the LOOKUP option in TCPIP.DATA to specify local searching before or instead of asking DNS.

2. Determine whether the name or address being queried is known to DNS if you expect to resolve the hostname using DNS. The following example looks for the name www.johndoe.com from the DNS at IP address 1.2.3.4:

   ```
   dig @1.2.3.4 www.johndoe.com -t any
   ```

   The command should return all resource records of any type from the DNS at 1.2.3.4 for www.johndoe.com. For more information about dig, see z/OS Communications Server: IP System Administrator's Commands.

3. If dig does not return all resource records, base your next course of action on the following conditions:

If dig. . .	Then . . .	Solution
Fails because it cannot contact DNS	You need to check your link to the DNS IP address.	See Chapter 4, "Diagnosing network connectivity problems," on page 25 to continue researching the problem.

If dig. . .	Then . . .	Solution
Fails because DNS reports that the resource was not found	www.johndoe.com is not a resource record known to DNS.	See the DNS administrator to add the name. As a temporary work-around, you might want to add the name to a local host file that the resolver searches. See z/OS Communications Server: IP Configuration Guide for information about local host files.
Succeeds	The problem in resolving the hostname using ping or another application might be in configuring the resolver, or might involve the contents of the resolver cache, or might result from the resolver treating the name server as unresponsive.	The **dig** command bypasses the resolver cache, search orders, local host files, and domain names that are appended by the resolver. The best way to check the configuration is to start the Trace Resolver. It is important to use the Trace Resolver in the environment where the application is failing because the application might be using a different TCPIP.DATA file, environment variables, or search order than the environment where the **dig** command was issued. The application might also be accessing the resolver cache, which might have inaccurate or outdated information. See "Steps for resolving caching problems" on page 872 for more cache information. The resolver might not be sending DNS queries generated by applications because the name server is considered to be unresponsive. See "Steps for responding to message EZZ9311E" on page 879 for more information.

4. If the Trace Resolver output indicates errors in the resolver configuration, issue the MODIFY RESOLVER,DISPLAY command and base your course of action on the following conditions:

If MODIFY RESOLVER,DISPLAY output...	Then...	Solution
Includes message EZD2039I WARNINGS ISSUED DURING RESOLVER INITIALIZATION	The resolver detected syntax errors or unrecognized setup statements when it was processing the resolver setup file statements during address space initialization. These errors might cause the resolver to use different settings than you intended.	Correct the errors in the resolver setup file. Issue the MODIFY RESOLVER,REFRESH,SETUP command to correct the resolver configuration.
Does not include message EZD2039I.	No syntax errors were in the resolver setup file. The configuration error might still be a value that is incorrectly specified on a resolver setup statement, or it might be an error in the application TCPIP.DATA file.	Correct the errors in the resolver setup file or application TCPIP.DATA file. Issue the MODIFY RESOLVER,REFRESH command to correct the configuration errors.

Results

You know that you are done when the application that previously failed to resolve the host name can now resolve it.

Steps for resolving caching problems

Determine whether a hostname contains inaccurate information in the resolver cache.

Before you begin

You need to know the exact hostname that is suspected to have inaccurate information in the resolver cache, and the environment in which the application was running (for example, TSO, UNIX, or batch).

Procedure

1. Determine whether the resolver cache contains any information about the host name, using the Netstat RESCache/-q report. The following command can be used to display information about the host name (which in this example is www.johndoe.com):

   ```
   netstat -q DETAIL -H www.john.doe
   ```

   The command should display all entries that exist in the resolver cache because of hostname-to-IP address resolution requests for www.johndoe.com. For more information about the Netstat RESCache/-q report, see z/OS Communications Server: IP System Administrator's Commands.
2. Determine, using the **dig** command, if the name or address being queried is known to the DNS name server if you expect to resolve the hostname using DNS. The following example looks for the name www.johndoe.com from the DNS at IP address 1.2.3.4:

   ```
   dig@1.2.3.4 www.johndoe.com -t any
   ```

   The command should return all resource records of any type from the DNS at 1.2.3.4 for www.johndoe.com. For more information about dig, see z/OS

Communications Server: IP System Administrator's Commands. The DNS name server to use for this search is the first name server listed in the NSINTERADDR list of name servers in the TCPIP.DATA data set used by the application experiencing the problems. For more information about NSINTERADDR, see z/OS Communications Server: IP Configuration Reference.

3. Base your next course of action on the results of the **dig** command:

If . . .	Then . . .	Solution
Dig fails because it cannot contact DNS	You need to check your link to the DNS IP address.	See Chapter 4, "Diagnosing network connectivity problems," on page 25 to continue researching the problem.
Dig fails because DNS reports that the resource was not found	www.johndoe.com is not a resource record known to DNS.	See the DNS administrator to add the name.
If dig succeeds, and the resolver cache has different information from what was returned by dig	The cache information might be information that was provided by a different name server, or might represent old information that was assigned a time-to-live (TTL) value that was excessive. The application might be acquiring inaccurate information about the hostname due to the resolver cache data.	Because the **dig** command bypasses the resolver cache, the dig command is unaffected by the resolver cache information. To remove the cache information, issue the MODIFY RESOLVER,FLUSH,ALL command to delete all entries from the cache. For more information about MODIFY FLUSH processing, see z/OS Communications Server: IP System Administrator's Commands.

If . . .	Then . . .	Solution
If dig succeeds, but the resolver cache has the same information as dig, or has no information about the hostname	The problem in resolving the hostname using ping or another application might be in configuring the resolver, or might result from the resolver treating the name server as unresponsive.	The **dig** command bypasses the resolver cache, search orders, local host files, and domain names that are appended by the resolver. The best way to check the configuration is to start the Trace Resolver. It is important to use the Trace Resolver in the environment where the application is failing because the application might be using a different TCPIP.DATA file, environment variables, or search order than the environment where the **dig** command was issued. The resolver might not be sending DNS queries that are generated by applications because the name server is considered to be unresponsive. See "Steps for responding to message EZZ9311E" on page 879 for more information.

4. If the Trace Resolver output indicates errors in the resolver configuration, issue the MODIFY RESOLVER,DISPLAY command and base your course of action on the following conditions:

If MODIFY RESOLVER,DISPLAY output . . .	Then . . .	Solution
Includes message EZD2039I WARNINGS ISSUED DURING RESOLVER INITIALIZATION	The resolver detected syntax errors or unrecognized setup statements when it was processing the resolver setup file statements during address space initialization. These errors might cause the resolver to use different settings than you intended.	Correct the errors in the resolver setup file. Issue the MODIFY RESOLVER,REFRESH,SETUP command to correct the resolver configuration.
Does not include message EZD2039I.	No syntax errors were in the resolver setup file. The configuration error might still be a value that is incorrectly specified on a resolver setup statement, or it might be an error in the application TCPIP.DATA file.	Correct the errors in the resolver setup file or application TCPIP.DATA file. Issue the MODIFY RESOLVER,REFRESH command to correct the configuration errors.

Results

You know that you are done when the application that previously failed to resolve the host name can now resolve it.

Steps for responding to message EZZ9308E

Message EZZ9308E is displayed when the resolver is monitoring the responsiveness of name servers, the network operator notification function is active, and the resolver determined that a name server is not responding to a significant percentage of queries over a 5-minute interval, where significant is defined by the setting of the UNRESPONSIVETHRESHOLD resolver setup statement. You might want to take action for your name servers when you see message EZZ9308E.

Procedure

Use the following steps to determine your course of action, if any, in response to message EZZ9308E:

1. Evaluate the scope of the problem with the name server.

   Answer the following questions to determine whether there is a significant issue with the name server:

Table 89. Message EZZ9308E name server resolution, part 1

Question	Implication
Is the monitoring function enabled when it should not be enabled?	The resolver might have detected syntax errors when processing the resolver setup file, and might have used the default setting for UNRESPONSIVETHRESHOLD.
Are there known network problems that are contributing to the situation?	If known network problems exist, then message EZZ9308E is likely a result of those problems, and the name server issues should clear up once the network problem is resolved.
Is only one name server unresponsive, or are multiple name servers unresponsive?	Resolver generates message EZZ9308E for each name server that is unresponsive. • If multiple name servers are reported as unresponsive at roughly the same time, it is likely a more systemic network problem than a problem with any given name server. • If only one name server is reported as unresponsive, it is more likely to be an issue with that specific name server.
Is this name server the primary, mission critical name server, or is it the backup name server?	Problems with the primary name server are likely to cause more disruption to the network and to your system, and thus are more likely to require intervention, than problems with a secondary name server.
Is the reported IP address not a valid address for contacting a name server?	One or more TCPIP.DATA data sets being used by applications on the systems have an incorrect IP address coded on an NSINTERADDR or NAMESERVER statement, and resolver is repeatedly attempting to send queries to that IP address.
Is the volume of requests that are failing significant?	Resolver reports the total number of resolver queries and total number of failures associated with this name server in message EZZ9310I, which is displayed when the name server is first identified as being unresponsive, and also at 5-minute intervals after that for as long as the name server remains unresponsive. High numbers of failures, or a significant percentage of a high number of requests, represents a larger disruption to your system than a small number of failures to a name server that is seldom used.

Table 89. Message EZZ9308E name server resolution, part 1 (continued)

Question	Implication
Is the name server responding, but not fast enough to be considered responsive by resolver?	Coding a very small RESOLVERTIMEOUT value might cause resolver to treat late arriving name server responses as failures to respond, even though the response from the name server is used by resolver to satisfy the API call. A less aggressive RESOLVERTIMEOUT value might alleviate the situation, or the situation could be ignored if slight network disruptions are thought to be the issue. Issue **dig** commands to query information at the name server being reported as unresponsive. Use the +time operand on the **dig** command to override the setting for RESOLVERTIMEOUT in order to determine whether the setting of RESOLVERTIMEOUT is possibly at fault.
Is the number of failures artificially high due to TCPIP.DATA settings?	The settings for RESOLVERUDPRETRIES and SEARCH might cause resolver to increment the failure count multiple times for a single resolver API call, thereby possibly exaggerating the impact of the failures to your system. • If multiple domain names are coded on the SEARCH statement, multiple searches for one hostname, with the unique domain names appended, might be attempted. • If a value greater than 1 is coded for RESOLVERUDPRETRIES, multiple attempts to contact the same name server for the same resource might be attempted. In either case, a small number of API calls could result in a significantly higher failure rate than would be expected for a lightly used name server. If these settings are combined with a small RESOLVERTIMEOUT value, the number could potentially be very high and yet the name server could be running normally.

2. Based on the answer to your evaluation of the scope of the problem, take the following actions:

Table 90. Message EZZ9308E name server resolution, part 2

Answer	Actions required
The problem is related to errors in the resolver setup file.	1. Identify and correct the errors in the resolver setup file. 2. Issue the MODIFY RESOLVER,REFRESH,SETUP=setup_filename command to correct the resolver configuration settings.
The problem is network-related.	1. Identify and correct the network problem. 2. Optionally, clear the eventual action messages from the operator console, or leave the message on the operator console and wait for the resolver to clear the message when the name server is again responsive.

Table 90. Message EZZ9308E name server resolution, part 2 (continued)

Answer	Actions required
The problem is related to a valid name server.	1. Identify and correct the problem with the name server. 2. Optionally, enable the autonomic quiescing of unresponsive name servers function so that the resolver automatically stops forwarding DNS queries generated by an application to the valid name server while it is unresponsive. The automatic quiescing function requires you to define a global TCPIP.DATA file. Review the information about the resolver and the global TCPIP.DATA file in the z/OS Communications Server: IP Configuration Guideto determine whether you can use a global TCPIP.DATA file. If you can use a global TCPIP.DATA file, enable the automatic quiescing function by performing the following steps: • If you do not have a resolver setup file, create one. • If you do not have a global TCPIP.DATA file, create one. Code one or more NSINTERADDR statements in the global TCPIP.DATA file, specifying the IP addresses of the name servers to be used in your environment. • Code the GLOBALTCPIPDATA statement in the resolver setup file, and specify the name of your global TCPIP.DATA on the statement. Code the UNRESPONSIVETHRESHOLD statement in the setup file, specifying a threshold percentage and also specifying the AUTOQUIESCE operand. Issue the MODIFY RESOLVER,REFRESH,SETUP=setup_filename command to cause the resolver to use the new threshold value and to automatically stop forwarding DNS queries generated by an application to unresponsive name servers. 3. Optionally, clear the eventual action messages from the operator console, or alternatively leave the message on the operator console and wait for the resolver to clear the message when the name server is again responsive.
The problem is related to an aggressive timeout value for resolver queries.	1. Identify the TCPIP.DATA data sets that have the small RESOLVERTIMEOUT. 2. Modify the TCPIP.DATA data sets to increase the timeout value. 3. Issue MODIFY RESOLVER,REFRESH to cause resolver to read the updated TCPIP.DATA data sets the next time any application using the data set issues a resolver request. 4. Optionally, clear the eventual action messages from the operator console, or alternatively leave the message on the operator console and wait for resolver to clear the message when the name server is again responsive.

Table 90. Message EZZ9308E name server resolution, part 2 (continued)

Answer	Actions required
The problem is related to an incorrect IP address.	1. Clear the eventual action messages from the operator console. 2. Identify the TCPIP.DATA data sets that have the incorrect IP address in the list of name servers to contact. 3. Modify the TCPIP.DATA data sets to eliminate the IP address. 4. Issue MODIFY RESOLVER,REFRESH to cause resolver to read the updated TCPIP.DATA data sets the next time any application using the data set issues a resolver request.
The problem is not considered to be an error, or is considered to be insignificant, or is a result of TCPIP.DATA settings that are adding to the problem.	1. Clear the eventual action messages from the operator console. 2. Optionally, if TCPIP.DATA settings are thought to be exaggerating the issue, and those settings can be modified without disrupting normal network processing, perform the following steps: • Identify the TCPIP.DATA data sets that have the settings that need to be changed. • Modify the TCPIP.DATA data sets to correct the RESOLVERUDPRETRIES or SEARCH settings. • Issue MODIFY RESOLVER,REFRESH to cause the resolver to use the new settings. 3. Optionally modify the value of threshold value for declaring a name server to be unresponsive in order to reduce the likelihood of future EZZ9308E messages being displayed for this name server. To modify the threshold value, perform the following steps: • If you do not have a resolver setup file, create one. • Code the desired percentage value for UNRESPONSIVETHRESHOLD in the setup file. Coding a value of zero will turn off the monitoring function. • Issue the MODIFY RESOLVER,REFRESH,SETUP=setup_filename command to cause the resolver to use the new threshold value.

3. Manually clearing message EZZ9308E does not mean that the name server is now responsive, so you might need to continue to monitor the state of the name server. Based on your actions in Step 2, do one of the following:

Table 91. Message EZZ9308E name server resolution, part 3

Monitoring function status	Actions required
You did not turn off the monitoring function.	Monitor the operator console for these resolver messages: 1. EZZ9310I, which the resolver issues to provide statistics for unresponsive name servers at 5-minute intervals. Use the statistics provided in EZZ9310I to ensure that the problem with this name server does not become a more serious problem, for instance that the percentage of failed messages to the name server does not become significantly higher. If the name server continues to fail to respond to messages for a longer period of time than expected, or the percentage of queries receiving no response increases significantly, re-evaluate the problem by using the questions in Step 1. 2. EZZ9309I, which the resolver issues to indicate that a previously unresponsive name server is now responding to queries at an acceptable level. One last instance of EZZ9310I is displayed with the EZZ9309I message. If EZZ9308E had not already been cleared by the operator, the resolver will clear the message from the console at this time.
You turned on the autonomic quiescing of unresponsive name servers function.	Issue the MODIFY RESOLVER,DISPLAY command to verify that the UNRESPONSIVETHRESHOLD statement is set to the percentage that you specified in the resolver setup file and that AUTOQUIESCE is set. If the name server remains unresponsive, the resolver will now issue message EZZ9311E, and will stop forwarding DNS queries generated by an application to the unresponsive name server.
You turned off the monitoring function.	Issue the MODIFY RESOLVER,DISPLAY command to verify that the UNRESPONSIVETHRESHOLD statement is set to 0. You will receive no new unresponsive messages, including EZZ9308E or EZZ9311E, until you restart the monitoring function.

Results

You know you are done when EZZ9308E no longer appears on the operator console and the name server is now responsive or the monitoring function is disabled or the autonomic quiescing function is active. If the monitoring function is still enabled and resolver subsequently determines that the name server is again unresponsive, a new message EZZ9308E is displayed.

Steps for responding to message EZZ9311E

You might want to take action regarding your name servers when you see message EZZ9311E.

About this task

Message EZZ9311E is displayed when the resolver is monitoring the responsiveness of name servers, the autonomic quiescing of unresponsive name servers function is active, and the resolver has determined that a name server is not responding to a significant percentage of queries over a 30-second interval where significant is defined by the setting of the UNRESPONSIVETHRESHOLD

resolver setup statement. The resolver will not forward DNS queries generated by an application to an unresponsive name server displayed in message EZZ9311E.

Guideline: If all possible name servers are unresponsive, the resolver sends queries generated by the application to the name servers rather than failing the application-query immediately.

Procedure

Use the following steps to determine your course of action, if any, in response to message EZZ9311E:

1. Evaluate the scope of the problem with the name server. Answer the following questions to determine whether there is a significant issue with the name server.

Table 92. Message EZZ9311E name server resolution, part 1

Question	Implication
Is the monitoring function enabled when it should not be enabled?	The resolver might have detected syntax errors when processing the resolver setup file, and might have used the default setting for UNRESPONSIVETHRESHOLD.
Are there known network problems that are contributing to the situation?	If known network problems exist, then message EZZ9311E is likely a result of those problems. The name server issues should stop when the network problem is resolved.
Is only one name server unresponsive, or are multiple name servers unresponsive?	The resolver generates message EZZ9311E for each name server that is unresponsive. • If multiple name servers are reported as unresponsive at roughly the same time, the network problem is probably a systemic problem, rather than a problem with any given name server. • If only one name server is reported as unresponsive, the problem is probably an issue with that specific name server.
Is this name server the primary, mission-critical name server, or is it the backup name server?	Problems with the primary name server are likely to cause more disruption to the network and to your system, and are more likely to require intervention than problems with a secondary name server.
Is the reported IP address a valid address for contacting a name server?	An incorrect IP address might be coded on an NSINTERADDR or NAMESERVER statement in your global TCPIP.DATA file. The resolver is repeatedly attempting to send DNS queries generated by an application or DNS polling queries to that IP address.
Is the volume of requests that are failing significant?	The resolver reports the total number of resolver queries and total number of failures associated with this name server in message EZZ9313I, which is displayed along with message EZZ9311E when the name server is identified as unresponsive. The resolver will not consider a name server to be unresponsive if less than 10 total DNS queries completed during the monitoring interval. A high percentage of failures to your primary name server represents a larger disruption to your system than a small number of failures to a name server that is seldom used.

Table 92. Message EZZ9311E name server resolution, part 1 (continued)

Question	Implication
Is the name server responding, but not fast enough to be considered responsive by the resolver?	Coding a low RESOLVERTIMEOUT value might cause the resolver to treat name server responses that arrive late as failures to respond, even though the response from the name server is used by the resolver to satisfy the API call. A higher RESOLVERTIMEOUT value might alleviate the situation. If you think that slight network disruptions are the problem, then you might choose to ignore the problem. Issue commands to query for information at the name server that is being reported as unresponsive. Use the +time operand on the command to override the setting for the RESOLVERTIMEOUT statement to determine whether the RESOLVERTIMEOUT setting is causing the problem.
Is the number of failures artificially high due to TCPIP.DATA settings?	The settings for the RESOLVERUDPRETRIES and SEARCH statements might cause the resolver to increment the failure count multiple times for a single resolver API call. Multiple increments can exaggerate the impact of the failures to your system. • If multiple domain names are coded on the SEARCH statement, the resolver might attempt multiple searches for one host name, with the unique domain names appended. • If a value greater than 1 is coded for RESOLVERUDPRETRIES, the resolver might attempt to contact the same name server for the same resource multiple times. In either case, a low number of API calls could result in a significantly higher failure rate than you would expect for a name server that is not used often. If these settings are combined with a low RESOLVERTIMEOUT value, the number of API calls could potentially be high and yet the name server could be running normally.

2. Based on the answer to your evaluation of the scope of the problem in Step 1, take the appropriate action listed in Table 93:

Table 93. Message EZZ9311E name server resolution, part 2

Answer	Actions required
The problem is related to errors in the resolver setup file.	1. Identify and correct the errors in the resolver setup file. 2. Issue the MODIFY RESOLVER,REFRESH,SETUP=setup_filename command to correct the resolver configuration settings.
The problem is network-related.	1. Identify and correct the network problem. 2. Optionally, clear the eventual action messages from the operator console, or leave the message on the operator console and wait for the resolver to clear the message when the name server is again responsive.

Table 93. Message EZZ9311E name server resolution, part 2 (continued)

Answer	Actions required
The problem is related to the name server.	1. Identify and correct the problem with the name server. 2. Optionally, clear the eventual action messages from the operator console, or leave the message on the operator console and wait for the resolver to clear the message when the name server is again responsive.
The problem is related to a low timeout value for resolver queries.	1. Modify the global TCPIP.DATA data set to increase the timeout value. 2. Issue the MODIFY RESOLVER,REFRESH command to cause the resolver to read the updated TCPIP.DATA data sets the next time any application using the data set issues a resolver request. 3. Optionally, clear the eventual action messages from the operator console, or leave the message on the operator console and wait for the resolver to clear the message when the name server is again responsive.
The problem is related to an incorrect IP address.	1. Clear the eventual action messages from the operator console. 2. Modify the global TCPIP.DATA data set to eliminate the IP address. 3. Issue the MODIFY RESOLVER,REFRESH command to cause the resolver to read the updated TCPIP.DATA data sets the next time any application issues a resolver request.
The problem is not considered to be an error or is a result of TCPIP.DATA settings that are adding to the problem.	1. Clear the eventual action messages from the operator console. 2. Optionally, if you think that TCPIP.DATA settings are exaggerating the issue and you can modify the settings without disrupting normal network processing, perform the following steps: a. Modify the global TCPIP.DATA data set to correct the RESOLVERUDPRETRIES or SEARCH settings. b. Issue the MODIFY RESOLVER,REFRESH command to cause the resolver to use the new settings. 3. Optionally, modify the threshold value for declaring a name server to be unresponsive to reduce the likelihood of future EZZ9311E messages being issued for this name server. To modify the threshold value, perform the following steps: a. Code the percentage value for the UNRESPONSIVETHRESHOLD statement in the setup file. The value 0 turns off the monitoring function. b. Issue the MODIFY RESOLVER,REFRESH,SETUP=setup_filename command to cause the resolver to use the new threshold value.

Table 93. Message EZZ9311E name server resolution, part 2 (continued)

Answer	Actions required
The problem is considered to be insignificant, or you no longer want the resolver to automatically stop forwarding DNS queries to unresponsive name servers.	Disable the autonomic quiescing of unresponsive name servers function. To disable the function, perform the following steps: 1. Remove the AUTOQUIESCE operand from the UNRESPONSIVETHRESHOLD statement in the resolver setup file. Code a percentage value for the UNRESPONSIVETHRESHOLD statement in the setup file. If you want to use the default value of 25%, you can delete the UNRESPONSIVETHRESHOLD statement from the setup file. The value 0 turns off the monitoring function. 2. Issue the MODIFY RESOLVER,REFRESH,SETUP=setup_filename command.

3. Manually clearing message EZZ9311E does not mean that the name server is now responsive, so you might need to continue to monitor the state of the name server. Based on your actions in Step 2, take the appropriate action in Table 94:

Table 94. Message EZZ9311E name server resolution, part 3

Monitoring function status	Actions required
You did not turn off the autonomic quiescing function.	1. Periodically issue the MODIFY RESOLVER,DISPLAY command. Use the failure rate information provided in the EZD2035I message associated with this name server to ensure that the problem with this name server does not become a more serious problem. For example, ensure that the percentage of failed DNS queries to the name server does not become significantly higher. If the name server continues to fail to respond to resolver DNS polling queries for a longer period of time than you expect, or the percentage of queries that receive no response increases significantly, re-evaluate the problem by going back to Step 1. 2. Monitor the operator console for message EZZ9312I, which the resolver issues to indicate that a previously unresponsive name server is responding to resolver DNS polling queries at an acceptable level. If EZZ9311E has not already been cleared by the operator, the resolver clears the message from the console.
You turned off the autonomic quiescing function.	Issue the MODIFY RESOLVER,DISPLAY command to verify that the UNRESPONSIVETHRESHOLD statement is set to the percentage that you specified in the resolver setup file (or the default value 25%), and to verify that AUTOQUIESCE is not displayed. If the name server remains unresponsive, the resolver issues message EZZ9308E and continues to forward DNS queries generated by an application to the name server.

Table 94. Message EZZ9311E name server resolution, part 3 (continued)

Monitoring function status	Actions required
You turned off the monitoring function completely.	Issue the MODIFY RESOLVER,DISPLAY command to verify that the UNRESPONSIVETHRESHOLD statement is set to 0. You will receive no new unresponsive messages, including messages EZZ9308E or EZZ9311E, until you restart the monitoring function.

Results

You know you are done when one of the following situations occurs:

- Message EZZ9311E no longer appears on the operator console and the name server is responsive
- The network operator notification function is active
- The monitoring function is disabled.

If the autonomic quiescing function is still enabled and resolver subsequently determines that the name server is again unresponsive, a new message EZZ9311E is displayed.

TRACE RESOLVER

You can activate Trace Resolver to collect Trace Resolver output. Trace Resolver output displays information about what the resolver looked for (the questions), where it looked (available name servers, resolver cache information, and local host table definitions), and what information was used to satisfy a resolver API request. Because the resolver runs in the address space of the application, you must collect Trace Resolver from the application address space. After the Trace Resolver output is collected, check the following information in the output:

- Fix or check any problems reported at the top of the trace. These are errors in the resolver data sets, or they might be errors detected in the resolver setup file during initialization of the resolver address space.
- The environment (z/OS UNIX or native MVS) that the application is running in determines where the resolver obtains the TCPIP.DATA statements to use. Are the data sets being used by the resolver the ones you expected? If not, see the search orders for data sets in z/OS Communications Server: IP Configuration Guide.
- Check that the expected MVS data sets or UNIX file system files are accessible by the user or batch job. Errors detected by a security product (for example, RACF) or OPEN services can generate messages that help indicate the problem. For example, IEC141I 013-C0 can be generated if a file does not have the correct permission bit settings to allow it to be read. RACF message ICH408I can be issued if no OMVS segment is defined or if insufficient authorization is granted to read a data set. See z/OS Communications Server: IP Configuration Guide for more information.
- Check the TCPIP.DATA parameter values, especially NAMESERVER, NOCACHE, NOCACHEREORDER, NSINTERADDR, NSPORTADDR and SEARCH. TCPIP.DATA parameters are explained in z/OS Communications Server: IP Configuration Reference.
- Check the questions posed by the Resolver to DNS or in searching the local host files. Are these the queries you expected?
- Check for the cache query and cache add attempts being attempted by the resolver as part of processing a resolver query. Are these the queries you

expected, and the results that you anticipated for those requests? Are the results from the cache-query attempts returning the information that you expected to be returned for those target resources? If you suspect that there are problems with the operations of the resolver cache, you will need to collect CTRACE records for further diagnosis. See "CTRACE - RESOLVER" on page 907 for more details.

- Look for errors or failures in the trace.
- Did DNS respond (if you expected it to)? If not, see whether DNS is active at the IP address you specified for NAMESERVER and NSINTERADDR and what port it is listening on. DNS logs can be helpful. Ask the DNS administrator for help.

  Tips: The resolver supports the Extension Mechanisms for DNS (EDNS0) standards, which permit DNS messages of greater than 512 bytes to be returned by DNS to the resolver; however, some network routers are configured to silently discard DNS messages of greater than 512 bytes. If the Trace Resolver suggests that DNS did not respond, verify that no routers between the resolver and DNS are discarding the messages.

  The resolver dynamically determines the EDNS0 capability of each DNS, based on responses or timeouts to DNS queries. The current view of the DNS capabilities is included in Trace Resolver output. You can issue the Netstat HOme/-h command to examine Trace Resolver output. Use the Trace Resolver information to verify that the resolver is using proper EDNS0 processing for a given DNS. Issue the MODIFY REFRESH command to force the resolver to dynamically relearn the EDNS0 capability of the name servers. Consider adjusting RESOLVERTIMEOUT values if timeout conditions cause the resolver to mistakenly avoid EDNS0 processing for a given DNS.
- The following are some common misunderstandings:
  - If the queried name server returns NXDOMAIN, the resolver does not continue to the next name server in the list. NXDOMAIN means the domain does not exist according to that name server.
  - The resolver only appends the specific names listed in the Search (or Domain) parameter. It does not attempt shorter versions of these. For example, if you look for "johndoe" and your search list has "anywhere.usa.com", the resolver looks for "johndoe.anywhere.usa.com" and "johndoe". (The order depends on the value of option ndots.) The Resolver does not look for "johndoe.anywhere" or "johndoe.anywhere.usa" or "johndoe.usa.com" or "johndoe.com".
  - The contents of any local hosts files are not cached in the system-wide resolver cache, but are saved separately for each task.
  - Negative cache entries are created to represent the following responses from a name server:
    - A response with a return code value of NXDOMAIN. This typically represents a host name that has no records of any type in the specified domain.
    - A response with a return code value of NOERROR when no answer records are returned. This represents a host name that does not have any records of the type that was requested (A, AAAA, etc.) in the specified domain, but does have some records of a different type.
    - A response with a return code value of NOERROR when the answer records returned represent canonical, or alias, names. This represents a resource that is officially known by other names in the specified domain, does not have any records of the type that was requested (A, AAAA, etc.), but does have some records for a different type of resource.

- You cannot use the cache reordering function to ensure different reordered resolution results for individual applications. The list of cached IP addresses is reordered in a round-robin manner for an individual cached host name regardless of which applications request resolution of the host name. Therefore, an individual application that issues multiple resolution requests for the same host name might receive the cached IP addresses in the same order in response to each resolution request, even if the cache reordering function is active. Resolver address sorting algorithms can also cause the same results to be generated in response to multiple requests for the same host name. See Cache reordering in z/OS Communications Server: IP Configuration Guide for more information.

Activate Trace Resolver output in one of the following ways:

- Specify the z/OS UNIX RESOLVER_TRACE environment variable or a SYSTCPT DD allocation. Specifying the RESOLVER_TRACE environment variable or allocating the SYSTCPT DDname dynamically activates Trace Resolver output regardless of the TCPIP.DATA or the _res structure resDebug specification. Dynamic activation of Trace Resolver can be useful when you are not sure where the TCPIP.DATA statements might be found.
- Specify the TCPIP.DATA statement TRACE RESOLVER or OPTIONS DEBUG. When using a TCPIP.DATA statement to activate the trace, have the trace activation statement as your very first statement. This ensures that the trace is in effect for all statements in the TCPIP.DATA specification.
- Set the debug option (resDebug) in an application _res structure.
- Enable the CTRACE TRACERES option to collect Trace Resolver output. You can use this option to debug resolver problems in one of the following conditions:
  - Long running started tasks or jobs cannot be restarted.
  - Long running started tasks or jobs have significant resolver API activity that often occurs under multiple tasks.

  You can use the CTRACE TRACERES option to collect the Trace Resolver output without restarting the started tasks or jobs. For more information, see "CTRACE - RESOLVER" on page 907.

The resolver uses the following search order to determine whether Trace Resolver output is necessary. The Trace Resolver data is contained in the specified output location. If the output location is not available for writing, the next search location is used. The default location for the Trace Resolver output in the z/OS UNIX environment is stdout. In the native MVS environment, it is as specified by the SYSPRINT DD.

1. The RESOLVER_TRACE environment variable (z/OS UNIX environment only).
2. The SYSTCPT DD allocation.
3. The TRACE RESOLVER or OPTIONS DEBUG statements. You must allocate STDOUT or SYSPRINT to generate trace data. The allocations need to exist in all operating environments including TSO, for example, your TSO Logon Procedure.
4. The resDebug bit set to on in the _res structure option field. STDOUT or SYSPRINT must be allocated or no trace data is generated.

Trace Resolver output can be written to any of the following:

- A TSO user terminal screen
- z/OS UNIX STDOUT
- JES SYSOUT

- An MVS Sequential data set (a member of a PDS is not supported). The data set must already exist or be allocated as new with the following DCB characteristics:
  - An LRECL between 80 and 256 with a RECFM of Fixed Block.
  - For an LRECL of 128 or larger, the last six print positions are the storage address of the MVS TCB that issued the resolver call. This can be helpful with multitask applications.
- A z/OS UNIX file system file. The file can either be an existing file or be dynamically allocated by the resolver when needed. The maximum line length used in the file is 255 characters. The last six print positions are the storage address of the MVS TCB that issued the resolver call. This can be helpful with multitask applications.

If the Trace Resolver output uses an MVS data set or z/OS UNIX file system file, the output is for the resolver services invoked by the last command or UNIX process. If possible, use SYSOUT=* or z/OS UNIX STDOUT to trace multiple resolver service invocations (for example, a multitask environment).

Specifying the Trace Resolver output location

Your environment determines the method to specify the Trace Resolver output location. This section includes the following environments:

- TSO
- z/OS UNIX
- MVS batch job
- z/OS UNIX batch

TSO environment

In the TSO environment, use one of the following to specify the Trace Resolver output location:

- For the user's terminal, enter the following:

  ```
  alloc dd(systcpt) da(*)
  ```

  When directing Trace Resolver output to a TSO terminal, define the screen size to be only 80 columns wide. Otherwise, trace output is difficult to read.
- For an existing MVS data set, enter the following:

  ```
  alloc dd(systcpt) da(appl.restrace)
  ```

  The user ID is used as the first qualifier for the data set. For example, if TSO USER1 entered the above command, user1 would be appended to the data set, as shown below:

  ```
  alloc dd(systcpt) da('user1.appl.restrace')
  ```

  To disable the Trace Resolver output, enter the following:

  ```
  free dd(systcpt)
  ```

z/OS UNIX shell environment

In the z/OS UNIX shell environment, use one of the following to specify the Trace Resolver output location:

- For STDOUT , enter the following:

  ```
  export RESOLVER_TRACE=STDOUT
  ```

  If needed, you can redirect STDOUT when the z/OS UNIX command is issued. If your application was compiled with the z/OS C/C++ Language Environment Native ASCII support or your application uses the FTP Client Application Programming Interface (API) do not use STDOUT. If you use STDOUT with ASCII programs the trace data is not readable. If you use STDOUT with the FTP

Client API the client's INIT command will fail. Instead send the trace data to an MVS data set or z/OS UNIX file system file as described below.

- For a new z/OS UNIX file system file or existing MVS data set, enter the following:

  ```
  export RESOLVER_TRACE=/tmp/myjob.resolv.trace
  export RESOLVER_TRACE="//appl.restrace"
  ```

  The user ID is used as the first qualifier for the data set. For example, if USER3 entered this command, user3 would be appended to the data set, as follows:

  ```
  export RESOLVER_TRACE="//'user3.appl.restrace'"
  ```

  To disable the Trace Resolver output, enter the following:

  ```
  set -A RESOLVER_TRACE
  ```

- For a z/OS UNIX file system file or an MVS data set that is already allocated to a ddname:

  ```
  export RESOLVER_TRACE="//dd:ddname"
  ```

  or

  ```
  export RESOLVER_TRACE="dd:ddname"
  ```

MVS batch job environment

In the MVS batch job environment, to use JES SYSOUT, enter the following statements:

```
//SYSTCPT  DD  SYSOUT=*
//SYSPRINT DD  SYSOUT=*
```

You must allocate either SYSTCPT or SYSPRINT DD if the TCPIP.DATA statements TRACE RESOLVER or OPTIONS DEBUG are specified. If neither is allocated, then no trace output is written.

z/OS UNIX batch environment

In the z/OS UNIX batch environment, use one of the following methods to specify the Trace Resolver output location:

- If the application resides in a z/OS UNIX file system file, use BPXBATSL to run the program. In this way, DD allocations is passed to the application. If the application does fork, the DD allocations are not passed to the new process, and the Trace Resolver output cannot be collected.
- To use the recommended JES SYSOUT, enter the following:

  ```
  //SYSTCPT  DD  SYSOUT=*
  ```

- Because STDOUT cannot be allocated to SYSOUT=* with BPXBATSL, use one of the following STDOUT DD JCL statements:

  ```
  //STDOUT   DD DISP=SHR,DSN=USER3.APPL.RESTRACE
  ```

  ```
  //STDOUT  DD PATH='/tmp/appl.stdout',
  //  PATHOPTS=(OWRONLY,OCREAT),
  //  PATHMODE=SIRWXU
  ```

  Note: In this example, OTRUNC is not specified on the PATHOPTS statement. This means the Trace Resolver output is appended to the z/OS UNIX file system file. To avoid z/OS UNIX file system full conditions, manually delete trace output that is no longer needed to ensure that the file does not fill the specified directory (for example, /tmp/).

  You must allocate either SYSTCPT or SYSPRINT DD if the TCPIP.DATA statements, TRACE RESOLVER or OPTIONS DEBUG, are specified. If neither is allocated, then no trace output is written.

- To pass the RESOLVER_TRACE environment variable using BPXBATSL or BPXBATCH, enter the following:

  ```
  //STDENV  DD JCL statement
  ```

  The following shows an example:

  ```
  //STDENV   DD  DISP=SHR,DSN=USER3.APPL.ENVIRON
  ```

  The STDENV data set can be a fixed or variable (nonspanned) record format type. It can contain multiple environment variables, as shown in the following sample:

  ```
  RESOLVER_TRACE=//'USER3.APPL.RESTRACE'
  _BPXK_SETIBMOPT_TRANSPORT=TCPCS
  ```

  Note:

  1. Environment variables must start in column 1, and the data set must not contain any sequence numbers because they would be treated as part of the environment variable.
  2. For the RESOLVER_TRACE environment variable, any blanks from a fixed-format STDENV data set is removed. Because this might not be true for all variables, a variable record format data set is recommended.
  3. For applications that fork, use of an MVS data set is recommended. If you use a z/OS UNIX file system file, a C03 ABEND might occur when the forked process ends.

Interpreting the Trace Resolver output

The following example shows the setup files that are used, the command that is used to invoke the trace, and the Trace Resolver output:

Setup files that are used for Trace Resolver:

Resolver Procedure:

```
//RESOLVER PROC PARMS='CTRACE(CTIRESFL)'
//*
//EZBREINI EXEC PGM=EZBREINI,REGION=0M,TIME=1440,PARM=&PARMS
//*
//SETUP   DD   DSN=TPOUSER.RESOLVER.SETUP.DATA,DISP=SHR,FREE=CLOSE
```

Setup File TPOUSER.RESOLVER.SETUP.DATA contains:

```
;
 DEFAULTTCPIPDATA('TPOUSER.RESOLVER.DEFAULT.DATA')
;
; GLOBALTCPIPDATA(/etc/tcpipglobal.data)
#
GLOBALTCPIPDATA('SYS1.TCPPARMS(RESGLOBL)')
```

Global TCPIP.DATA file SYS1.TCPPARMS(RESGLOBL) contains:

```
# Note that DOMAIN is ignored because SEARCH is mutually exclusive
# and SEARCH appears after DOMAIN.
Domain abcxyz
; Note that SEARCH can be specified on multiple lines.
 SEARCH tcp.raleigh.ibm.com raleigh.ibm.com
 SEARch ibm.com  com  uk
 SEARch gov
 1a Search mil
SORTLIST 0.0.19.0/0.0.255.0 0.0.18.99/0.0.255.255 0.42.17.0/0.255.255.0
SORTLIST 129.42.16.0/255.255.255.0
 1b Sortlist 9.0.0.0
NSinterAddr  111.111.111.111
NameServer   9.67.128.255 ; not a server
NSinterAddr 2001::9:43:80:52
NSportAddr 53
 2 ;ResolveVia UDP
ResolverTimeout 3
```

```
ResolverUdpRetries 1  ; 1 means 1 total try, 0 would be no tries at all
 1c loaddbcstables unknown
loaddbcstables big5
 3   MVS000:  Hostname MVS026
```

Default TCPIP.DATA file TPOUSER.RESOLVER.DEFAULT.DATA contains:

```
; TRACE RESOLVER
DatasetPrefix USER1
TcpipJobname TCPCS3
Hostname VIC097
; trace c sockets
; alwayswto no
; messagecase whoknows
; loaddbcstables tbd
```

Note: For this example, this file exists but is not used to obtain this example Trace Resolver output.

Local TCPIP.DATA file USER55.TCPIP.DATA contains:

```
   ; trace resolver
   DATASETPREFIX USER55
   # If an option is coded multiple times but can only have 1 value,
   # the last occurrence is used.
   TCPIPjobname TCPCS2
   TCPIPjobname TCPCS
 3   HostName MVS000
   DomainOrigin edu
   ;
   NameServer   127.0.0.1  ; loopback
   #
 2   ResolveVia TCP
   ResolverTimeout 22
 1d  alwayswto xyz
   messagecase mixed
   loaddbcstables schinese
```

TSO commands issued to obtain the trace (gethostbyname):

```
alloc dd(systcpt) dsn(traceres) reuse
 4 invoke a REXX application which issues gethostbyname for TESTBEN46.SVT390.COM
```

Trace Resolver output in USER55.TRACERES contains (gethostbyname):

```
 5 Resolver Trace Initialization Complete -> 2009/06/18 15:51:51.733708
 1a res_init Skipped option(s) on line     8: SYS1.TCPPARMS(RESGLOBL)
 1b res_init Skipped option(s) on line    11: SYS1.TCPPARMS(RESGLOBL)
 1c res_init Parse error on line    19: SYS1.TCPPARMS(RESGLOBL)
 1d res_init Parse error on line    14: USER55.TCPIP.DATA

 6  res_init Resolver values:
 35  Setup file warning messages = No
 36  CTRACE TRACERES option = No
  Global Tcp/Ip Dataset  = SYS1.TCPPARMS(RESGLOBL)
  Default Tcp/Ip Dataset = TPOUSER.RESOLVER.DEFAULT.DATA
  Local Tcp/Ip Dataset   = USER55.TCPIP.DATA
  Translation Table      = Default
  UserId/JobName         = USER55
   19 Caller API             = TCP/IP Rexx Sockets
   20 Caller Mode            = EBCDIC
   32 System Name            = MVS000 (from VMCF)
   34 UnresponsiveThreshold  = 25
  (L) DataSetPrefix = USER55
   3 (G) HostName      = MVS026
  (L) TcpIpJobName  = TCPCS
   3 (G) Search        = tcp.raleigh.ibm.com
                       raleigh.ibm.com
                       ibm.com
```

```
                      com
                      uk
                      gov
    (G) SortList      = 0.0.19.0/0.0.255.0
                      0.0.18.99/0.0.255.255
                      0.42.17.0/0.255.255.0
                      129.42.16.0/255.255.255.0
     3 (G) NameServer    = 111.111.111.111
                        EDNS0 Support = Unknown
                      9.67.128.255
                        EDNS0 Support = Unknown
                      2001::9:43:80:52
                        EDNS0 Support = Unknown
    (G) NsPortAddr    = 53            (G) ResolverTimeout    = 3
     2 (*) ResolveVia    = UDP           (G) ResolverUdpRetries = 1
    (*) Options NDots = 1
    (*) SockNoTestStor
    (*) AlwaysWto     = NO            (L) MessageCase        = MIXED
     11 (*) LookUp        = DNS LOCAL
|    25 (*) Cache
|    37 (*) CacheReorder
    (G) LoadDbcsTable = BIG5
   res_init Succeeded
   res_init Started: 2009/06/18 15:51:51.800668
   res_init Ended: 2009/06/18 15:51:51.800676
   ********************************************************************************
    4 GetHostByName Started: 2009/06/18 15:51:51.815838
   GetHostByName Resolving Name:  TESTBEN46.SVT390.COM
   GetHostByName Stack Name:   TCPCS
   res_search(TESTBEN46.SVT390.COM, C_IN, T_A)
   res_search Host Alias Search found no alias
   res_querydomain(TESTBEN46.SVT390.COM., , C_IN, T_A)
   res_querydomain resolving name:  TESTBEN46.SVT390.COM.
   res_query(TESTBEN46.SVT390.COM., C_IN, T_A)
    26 Querying resolver cache for TESTBEN46.SVT390.COM.
    EZBRECFR:  RetVal = 0, RC = 0, Reason = 0x00000000
     26a No cache information was available
   res_mkquery(QUERY, TESTBEN46.SVT390.COM., C_IN, T_A)
    7 res_mkquery created message:
   * * * * * Beginning of Message * * * * *
    Query Id:                18693
    Flags:                   00000001 00000000
    Flags set:               recurDes
    OpCode:                  QUERY
    Response Code:           NOERROR

    Number of Question RRs:  1
    Question 1:
    TESTBEN46.SVT390.COM
    Type (0X0001) T_A  Class (0X0001) C_IN

    Number of Answer RRs:  0
    Number of Authority RRs:  0
     7a Number of Additional RRs:  1
    Additional 1:
    Type (0X0029) T_OPT  UDP Payload (0X0C00) 3072
    Extended RCODE 0  Version 0  Flags 0000
   * * * * * End of Message * * * * *
    21 res_send Name Server Capabilities
    21b  Monitoring intervals used = 5
   Name server 111.111.111.111
      EDNS0 Support = unknown
      Queries sent  = 0
      Failures      = 0
      Percentage    = 0%
    21a  Use this name server for queries
   Name server 9.67.128.255
```

```
   EDNS0 Support = unknown
   Queries sent  = 0
   Failures      = 0
   Percentage    = 0%
   Use this name server for queries
 Name server 2001::9:43:80:52
   EDNS0 Support = unknown
   Queries sent  = 0
   Failures      = 0
   Percentage    = 0%
   Use this name server for queries
 8 res_send sending query to Name Server 111.111.111.111
 27 DNS Communication Started: 2009/06/18 15:51:51.816987
 31 BPX1SOC:  RetVal = 0, RC = 0, Reason = 0x00000000, Type=IPv6
 22 No OPT RR record sent on request to 111.111.111.111
 BPX1STO:  RetVal = 38, RC = 0, Reason = 0x00000000
 BPX1AIO Sched:  RetVal = 1, RC = 0, Reason = 0x00000000
 BPX1AIO RECVMSG :  From 111.111.111.111
                    RetVal=502, RC=0, Reason=0x00000000
 UDP Data Length: 502
 9 res_send received data via UDP.  Message received:
* * * * * Beginning of Message * * * * *
 Query Id:                18693
 Flags:                   10000111 10000000
 Flags set:               resp auth trunc recurDes recurAvl
 OpCode:                  QUERY
 Response Code:           NOERROR

 Number of Question RRs:  1
 Question 1:
 TESTBEN46.SVT390.COM
 Type (0X0001) T_A  Class (0X0001) C_IN

 Trace terminated due to truncation condition
* * * * * End of Message * * * * *
 27a DNS Communication Ended: 2009/06/18 15:51:51.836189 time used 00:00:00.019202
 8 res_send Sending query to Name Server 111.111.111.111
 27 DNS Communication Started: 2009/06/18 15:51:51.836812
 22a EDNS0 Probe request sent to 111.111.111.111 id=18694
 BPX1STO:  RetVal = 49, RC = 0, Reason = 0x00000000
 BPX1AIO Sched:  RetVal = 1, RC = 0, Reason = 0x00000000
 BPX1AIO RECVMSG :  From 111.111.111.111
                    RetVal=553, RC=0, Reason=0x00000000
 UDP Data Length: 553
 9a res_send received data via UDP.  Message received:
* * * * * Beginning of Message * * * * *
 Query Id:                18694
 Flags:                   10000101 10000000
 Flags set:               resp auth recurDes recurAvl
 OpCode:                  QUERY
 Response Code:           NOERROR

 Number of Question RRs:  1
 Question 1:
 TESTBEN46.SVT390.COM
 Type (0X0001) T_A  Class (0X0001) C_IN

 Number of Answer RRs:  29
 Answer 1:
 TESTBEN46.SVT390.COM
 Type (0X0001) T_A  Class (0X0001) C_IN
  30 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 197.14.103.5
 Answer 2:
 TESTBEN46.SVT390.COM
 Type (0X0001) T_A  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
```

```
197.14.103.6
Answer 3:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.7
Answer 4:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.8
Answer 5:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.9
Answer 6:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.10
Answer 7:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.11
Answer 8:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.12
Answer 9:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.13
Answer 10:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.14
Answer 11:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.15
Answer 12:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.16
Answer 13:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.17
Answer 14:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.18
Answer 15:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.19
Answer 16:
```

```
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.20
Answer 17:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.21
Answer 18:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.22
Answer 19:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.23
Answer 20:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.24
Answer 21:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.25
Answer 22:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.26
Answer 23:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.27
Answer 24:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.28
Answer 25:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.29
Answer 26:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.1
Answer 27:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.2
Answer 28:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
197.14.103.3
Answer 29:
TESTBEN46.SVT390.COM
Type (0X0001) T_A  Class (0X0001) C_IN
```

```
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 197.14.103.4

 Number of Authority RRs:  1
 Authority 1:
 SVT390.COM
 Type (0X0002) T_NS  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 sdistcici.SVT390.COM

  24 Number of Additional RRs:  2
 Additional 1:
 sdistcici.SVT390.COM
 Type (0X0001) T_A  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 197.11.235.51
 Additional 2:
 Type (0X0029) T_OPT  UDP Payload (0X1000) 4096
 Extended RCODE 0  Version 0  Flags 0000
* * * * * End of Message * * * * *
  23 Oversized reply to probe request from 111.111.111.111
 BPX1CLO:  RetVal = 0, RC = 0, Reason = 0x00000000
 27a DNS Communication Ended: 2009/06/18 15:51:51.841691 time used 00:00:00.004879
 33   Name Server Capability Updates
 Name server 111.111.111.111
   EDNS0 Support = up-level
   Queries sent  = 1
   Failures      = 0
res_send Succeeded
  28 Attempting to cache results for TESTBEN46.SVT390.COM.
 EZBRECAR:  RetVal = 0, RC = 0, Reason = 0x00000000
 Cache information was saved
res_query Succeeded
res_querydomain Succeeded
res_search Succeeded
 10 GetHostByName Succeeded:  IP Address(es) found:
  IP Address(1) is 197.14.103.5
  IP Address(2) is 197.14.103.6
  IP Address(3) is 197.14.103.7
  IP Address(4) is 197.14.103.8
  IP Address(5) is 197.14.103.9
  IP Address(6) is 197.14.103.10
  IP Address(7) is 197.14.103.11
  IP Address(8) is 197.14.103.12
  IP Address(9) is 197.14.103.13
  IP Address(10) is 197.14.103.14
  IP Address(11) is 197.14.103.15
  IP Address(12) is 197.14.103.16
  IP Address(13) is 197.14.103.17
  IP Address(14) is 197.14.103.18
  IP Address(15) is 197.14.103.19
  IP Address(16) is 197.14.103.20
  IP Address(17) is 197.14.103.21
  IP Address(18) is 197.14.103.22
  IP Address(19) is 197.14.103.23
  IP Address(20) is 197.14.103.24
  IP Address(21) is 197.14.103.25
  IP Address(22) is 197.14.103.26
  IP Address(23) is 197.14.103.27
  IP Address(24) is 197.14.103.28
  IP Address(25) is 197.14.103.29
  IP Address(26) is 197.14.103.1
  IP Address(27) is 197.14.103.2
  IP Address(28) is 197.14.103.3
  IP Address(29) is 197.14.103.4
GetHostByName Ended: 2009/06/18 15:51:51.846582
********************************************************************************
```

TSO commands issued to obtain the trace (getaddrinfo):

```
alloc dd(systcpt) dsn(traceres) reuse
12 CS V1R12: Pinging host TESTBEN46.SVT390.COM
at IPv6 address 2000:197:14:103::15
Ping #1 timed out
```

Trace Resolver output in USER55.TRACERES contains (getaddrinfo):

```
 5 Resolver Trace Initialization Complete -> 2009/06/18 15:56:52.840756
 1a res_init Skipped option(s) on line     8: SYS1.TCPPARMS(RESGLOBL)
 1b res_init Skipped option(s) on line    11: SYS1.TCPPARMS(RESGLOBL)
 1c res_init Parse error on line    19: SYS1.TCPPARMS(RESGLOBL)
 1d res_init Parse error on line    14: USER55.TCPIP.DATA

 6 res_init Resolver values:
 35  Setup file warning messages = No
 36  CTRACE TRACERES option = No
 Global Tcp/Ip Dataset  = SYS1.TCPPARMS(RESGLOBL)
 Default Tcp/Ip Dataset = TPOUSER.RESOLVER.DEFAULT.DATA
 Local Tcp/Ip Dataset   = USER55.TCPIP.DATA
 Translation Table      = Default
 UserId/JobName         = USER55
 19 Caller API             = TCP/IP Sockets Extended
 20 Caller Mode            = EBCDIC
 32 System Name            = MVS000 (from VMCF)
 34 UnresponsiveThreshold  = 25
(L) DataSetPrefix = USER55
 3 (G) HostName      = MVS026
 (L) TcpIpJobName  = TCPCS
 3 (G) Search        = tcp.raleigh.ibm.com
                     raleigh.ibm.com
                     ibm.com
                     com
                     uk
                     gov
 (G) SortList      = 0.0.19.0/0.0.255.0
                     0.0.18.99/0.0.255.255
                     0.42.17.0/0.255.255.0
                     129.42.16.0/255.255.255.0
 3 (G) NameServer    = 111.111.111.111
                       EDNS0 Support = Up-level
                     9.67.128.255
                       EDNS0 Support = Unknown
                     2001::9:43:80:52
                       EDNS0 Support = Unknown
 (G) NsPortAddr    = 53            (G) ResolverTimeout    = 3
 2 (*) ResolveVia    = UDP           (G) ResolverUdpRetries = 1
 (*) Options NDots = 1
 (*) SockNoTestStor
 (*) AlwaysWto     = NO            (L) MessageCase        = MIXED
 11 (*) LookUp        = DNS LOCAL
 25 (*) Cache
 37 (*) CacheReorder
 (G) LoadDbcsTable = BIG5
res_init Succeeded
res_init Started: 2009/06/18 15:56:52.894280
res_init Ended: 2009/06/18 15:56:52.894288
********************************************************************************
res_init Started: 2009/06/18 15:56:52.900190
res_init Ended: 2009/06/18 15:56:52.900197
********************************************************************************
 12 GetAddrInfo Started: 2009/06/18 15:56:52.900968
GetAddrinfo Invoked with following inputs:
   Host Name:  TESTBEN46.SVT390.COM
   No Service operand specified
   Hints parameter supplied with settings:
       ai_family = 0, ai_flags = 0x00000062
```

```
      ai_protocol = 0, ai_socktype = 0
 13 GetAddrInfo Opening Socket for IOCTLs
  31 BPX1SOC:  RetVal = 0, RC = 0, Reason = 0x00000000, Type=IPv6
 BPX1IOC:  RetVal = 0, RC = 0, Reason = 0x00000000
GetAddrInfo Opened Socket 0x00000000
 14 GetAddrInfo Both IPv4 and IPv6 Interfaces Exist
GetAddrInfo Host Alias Search found no alias
res_querydomain(TESTBEN46.SVT390.COM., , C_IN, T_AAAA)
res_querydomain resolving name:  TESTBEN46.SVT390.COM.
 15 res_query(TESTBEN46.SVT390.COM., C_IN, T_AAAA)
  26 Querying resolver cache for TESTBEN46.SVT390.COM.
 EZBRECFR:  RetVal = 0, RC = 0, Reason = 0x00000000
  26b No cache information was available
res_mkquery(QUERY, TESTBEN46.SVT390.COM., C_IN, T_AAAA)
 7 res_mkquery created message:
* * * * * Beginning of Message * * * * *
 Query Id:                52138
 Flags:                   00000001 00000000
 Flags set:               recurDes
 OpCode:                  QUERY
 Response Code:           NOERROR

 Number of Question RRs:  1
 Question 1:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN

 Number of Answer RRs:  0
 Number of Authority RRs:  0
 7a Number of Additional RRs:  1
 Additional 1:
 Type (0X0029) T_OPT  UDP Payload (0X0C00) 3072
 Extended RCODE 0  Version 0  Flags 0000
* * * * * End of Message * * * * *
 21 res_send Name Server Capabilities
 21b Monitoring intervals used = 5
 Name server 111.111.111.111
   EDNS0 Support = up-level
   Queries sent  = 1
   Failures      = 0
   Percentage    = 0%
  21a Use this name server for queries
Name server 9.67.128.255
   EDNS0 Support = unknown
   Queries sent  = 0
   Failures      = 0
   Percentage    = 0%
   Use this name server for queries
 Name server 2001::9:43:80:52
   EDNS0 Support = unknown
   Queries sent  = 0
   Failures      = 0
   Percentage    = 0%
   Use this name server for queries
 8 res_send Sending query to Name Server 111.111.111.111
 27 DNS Communication Started: 2009/06/18 15:56:52.902525
 22b OPT RR record included on request to 111.111.111.111
 BPX1STO:  RetVal = 49, RC = 0, Reason = 0x00000000
 BPX1AIO Sched:  RetVal = 1, RC = 0, Reason = 0x00000000
 BPX1AIO RECVMSG :  From 111.111.111.111
                    RetVal=649, RC=0, Reason=0x00000000
 UDP Data Length: 649
 9b res_send received data via UDP.  Message received:
* * * * * Beginning of Message * * * * *
 Query Id:                52138
 Flags:                   10000101 10000000
 Flags set:               resp auth recurDes recurAvl
```

```
OpCode:                   QUERY
Response Code:            NOERROR

Number of Question RRs:  1
Question 1:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN

Number of Answer RRs:  20
Answer 1:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
 30 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::18
Answer 2:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::19
Answer 3:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::20
Answer 4:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::1
Answer 5:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::2
Answer 6:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::3
Answer 7:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::4
Answer 8:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::5
Answer 9:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::6
Answer 10:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::7
Answer 11:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
2000:197:14:103::8
Answer 12:
TESTBEN46.SVT390.COM
Type (0X001C) T_AAAA  Class (0X0001) C_IN
```

```
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::9
 Answer 13:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::10
 Answer 14:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::11
 Answer 15:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::12
 Answer 16:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::13
 Answer 17:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::14
 Answer 18:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::15
 Answer 19:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::16
 Answer 20:
 TESTBEN46.SVT390.COM
 Type (0X001C) T_AAAA  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 2000:197:14:103::17

 Number of Authority RRs:  1
 Authority 1:
 SVT390.COM
 Type (0X0002) T_NS  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 sdistcici.SVT390.COM

  24 Number of Additional RRs:  2
 Additional 1:
 sdistcici.SVT390.COM
 Type (0X0001) T_A  Class (0X0001) C_IN
 TTL:  86400 (1 days, 0 hours, 0 minutes, 0 seconds)
 197.11.235.51
 Additional 2:
 Type (0X0029) T_OPT  UDP Payload (0X1000) 4096
 Extended RCODE 0  Version 0  Flags 0000
* * * * * End of Message * * * * *
 27a DNS Communication Ended: 2009/06/18 15:56:52.908818 time used 00:00:00.006293
 33a   Name Server Capability Updates
 Name server 111.111.111.111
   Queries sent  = 1
   Failures      = 0
res_send Succeeded
  28 Attempting to cache results for TESTBEN46.SVT390.COM.
```

```
 EZBRECAR:  RetVal = 0, RC = 0, Reason = 0x00000000
 Cache information was saved
res_query Succeeded
res_querydomain Succeeded
res_querydomain(TESTBEN46.SVT390.COM., , C_IN, T_A)
res_querydomain resolving name:  TESTBEN46.SVT390.COM.
 16 res_query(TESTBEN46.SVT390.COM., C_IN, T_A)
  26 Querying resolver cache for TESTBEN46.SVT390.COM.
 EZBRECFR:  RetVal = 0, RC = 0, Reason = 0x00000000
  26c Cache data from 111.111.111.111 was retrieved
res_query Succeeded
res_querydomain Succeeded
 17 GetAddrInfo Returning Zero as Port Number
GetAddrInfo Built 35 Addrinfos
 13a GetAddrInfo Closing IOCTL Socket 0x00000000
 BPX1CLO:  RetVal = 0, RC = 0, Reason = 0x00000000
 18 GetAddrInfo Succeeded:  IP Address(es) found:
  IP Address(1) is 2000:197:14:103::15
  IP Address(2) is 2000:197:14:103::14
  IP Address(3) is 2000:197:14:103::17
  IP Address(4) is 2000:197:14:103::16
  IP Address(5) is 2000:197:14:103::11
  IP Address(6) is 2000:197:14:103::10
  IP Address(7) is 2000:197:14:103::13
  IP Address(8) is 2000:197:14:103::12
  IP Address(9) is 2000:197:14:103::19
  IP Address(10) is 2000:197:14:103::18
  IP Address(11) is 2000:197:14:103::5
  IP Address(12) is 2000:197:14:103::4
  IP Address(13) is 2000:197:14:103::7
  IP Address(14) is 2000:197:14:103::6
  IP Address(15) is 2000:197:14:103::1
  IP Address(16) is 2000:197:14:103::3
  IP Address(17) is 2000:197:14:103::2
  IP Address(18) is 2000:197:14:103::9
  IP Address(19) is 2000:197:14:103::8
  IP Address(20) is 2000:197:14:103::20
  IP Address(21) is 197.14.103.5
  IP Address(22) is 197.14.103.6
  IP Address(23) is 197.14.103.7
  IP Address(24) is 197.14.103.8
  IP Address(25) is 197.14.103.9
  IP Address(26) is 197.14.103.10
  IP Address(27) is 197.14.103.11
  IP Address(28) is 197.14.103.12
  IP Address(29) is 197.14.103.13
  IP Address(30) is 197.14.103.14
  IP Address(31) is 197.14.103.15
  IP Address(32) is 197.14.103.16
  IP Address(33) is 197.14.103.17
  IP Address(34) is 197.14.103.18
  IP Address(35) is 197.14.103.19
 29 GetAddrInfo Ended: 2009/06/18 15:56:52.914971
********************************************************************************
FreeAddrInfo Started: 2009/06/18 15:56:52.917611
FreeAddrInfo Called to free addrinfo structures
FreeAddrInfo Succeeded, Freed 35 Addrinfos
FreeAddrInfo Ended: 2009/06/18 15:56:52.918496
********************************************************************************
```

The following section describes highlighted numbered areas of the example setup files and example Trace Resolver output.

1

Errors that are deliberately entered into this example to show action taken:

a

Line 8 in the global file specifies seven SEARCH values; the maximum number that is allowed is 6. The seventh value is ignored.

b

Line 11 in the global file specifies five SORTLIST values; the maximum number that is allowed is 4. The fifth value is ignored.

c

Line 19 in the global file has a value for LOADDBCSTABLES that is not valid. The value is ignored.

d

Line 14 in the local file has a value for ALWAYSWTO that is not valid. The value is ignored and the default is used.

2

The ResolveVia field specifies UDP even though the local file indicated RESOLVEVIA TCP. UDP is used because GLOBALTCPIPDATA is being used. If a global file is used, then all resolver-related TCPIP.DATA statements must be specified in it. If the resolver statements are not specified, then default values are assigned. In this example, resolver statements are not specified as shown by RESOLVEVIA being commented out in the global file.

3

A local file cannot override the global file for any value. The global file specifies the hostname, therefore the local file value of MVS000 does not override the global value of MVS026. Likewise, since there is a GLOBALTCPIPDATA specified all resolver-related statements in a local file are ignored (for example, DOMAINORIGIN, NAMESERVER and RESOLVERTIMEOUT).

The list of name servers to be queried also indicates whether the resolver considers the name server to support the Extension Mechanisms for DNS (EDNS0) capability. The resolver might believe that the name server supports EDNS0 (EDNS0 Support = Up-level), that the name server does not support EDNS0 (EDNS0 Support = Down-level), or that the capability of the name server is undetermined (EDNS0 Support = Unknown).

4

A REXX application calls GetHostByName at the indicated local date and time. The flow through the resolver API calls shows the parameters that are being passed.

5

Trace output reports the date and time the TCPIP.DATA statements were processed.

6

The `res_init()` resolver initialization values are reported. These values are the values actually being used by the resolver, with an indication of the origin of the value. The indicators are:

***** Default value

A Modified by application

D Default file (not used if the local file is found)

E Environment variable

G Global file

L Local file

7

`res_mkquery` creates a DNS message (from Beginning of Message to End of Message). The message is interpreted, and flags and codes are spelled out.

a `res_mkquery` appends an additional RR record (the OPT RR) to the request that is being built. The presence of the OPT RR record on the request indicates that the resolver supports the EDNS0 function and that UDP packets of up to 3072 bytes can be sent to the resolver (instead of the normal 512-byte limitation).

8

`res_send` sends the query to the name server. The res_send function calls several z/OS UNIX functions; the indentation of the lines that follow res_send indicate res_send was the caller. The IP address of the DNS is also displayed.

9

`res_send` receives a message from DNS. This message is truncated because the amount of data that the DNS has available to send regarding the resource is greater than 512 bytes.

a `res_send` receives a message from DNS. This time, because the DNS query from resolver contained the OPT RR record that indicates up to 3072 bytes of UDP data could be returned, the DNS returns all 553 bytes of data that use UDP protocols. The total number of Answer records in this response is 29.

b `res_send` receives a message from DNS. This message is also greater than 512 bytes, but since the OPT RR record is included on the query, the full amount of data can be sent on the first UDP response from the name server. The total number of Answer records in this response is 20.

10

The GetHostByName function reports success and lists the IP addresses returned. If addresses matched any of the values in the SORTLIST definitions, the order of the addresses would have been modified to match the SORTLIST specification.

11

LookUp specifies the order in which the DNS and the local host file are to be used for name resolution. There are four possible search orders:

LOOKUP DNS LOCAL (DNS search first)

LOOKUP LOCAL DNS (local host file search first)

LOOKUP DNS (only DNS search)

LOOKUP LOCAL (only local host file search)

12

Ping calls GetAddrinfo at the indicated local date and time. The flow through the resolver API calls shows the parameters that are being passed.

ai_family = 0 means that AF_UNSPEC is specified

ai_flags = x'00000062' means that AI_CANNONNAMEOK, AI_ALL, and AI_ADDRCONFIG are specified

ai_protocol = 0 and ai_socktype = 0 means that protocol and socktype are not specified

For more information about input values of getaddrinfo, see z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

13

To honor the setting of ai_ADDRCONFIG, the Resolver must query the stacks to determine whether IPv6 or IPv4 interfaces exist (the results of the query are shown in message **14**). A socket, separate from the one used to send DNS queries, is opened for communicating with the stacks.

a The socket that is used for communicating with the stacks is closed before finishing Getaddrinfo processing.

14

The resolver detected that the system can handle both IPv4 and IPv6 addresses.

15

Because the system can handle both IPv4 and IPv6, and ai_ALL is specified, the resolver sends the IPv6 query (T_AAAA) for IPv6 to DNS first. For an explanation of how resolver decides to send an IPv6 or IPv4 query to DNS, see z/OS Communications Server: IPv6 Network and Application Design Guide.

16

The resolver prepares to send the IPv4 query (T_A). For an explanation of how resolver decides to send an IPv6 or IPv4 query to DNS, see z/OS Communications Server: IPv6 Network and Application Design Guide.

17

Because no Service operand is passed as input to Getaddrinfo, there is no service resolution to perform, so any sockaddr returned has a port number=0.

18

Before the resolver returns resolved addresses to the application, the resolver sorts all addresses so that the most preferable is the first in the address chain. For more information, see z/OS Communications Server: IPv6 Network and Application Design Guide.

19

The caller API value indicates which search order is used by the resolver for any required local table usage. The following caller API values indicate that the z/OS UNIX environment search order is used:

1. Language Environment C Sockets

   For more information about z/OS XL C/C++ runtime library functions, see z/OS XL C/C++ Runtime Library Reference.

2. Unix System Services

   For more information about z/OS UNIX System Services (z/OS UNIX) callable services, see z/OS UNIX System Services Programming: Assembler Callable Services Reference.

The following caller API values indicate that the native MVS environment search order is used:

1. TCP/IP Pascal Sockets

   For more information about z/OS Communications Server PASCAL API, see z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

2. TCP/IP C Sockets

For more information about z/OS Communications Server C/C++ API, see z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

3. TCP/IP Rexx Sockets

   For more information about z/OS Communications Server REXX API, see z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

4. TCP/IP Sockets Extended

   For more information about z/OS Communications Server Callable and Macro API, see z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference.

20

The Caller Mode value indicates the representation of any input characters as being either in EBCDIC or ASCII.

21

Because the resolver awareness of the EDNS0 capability of name servers is maintained on a system-wide level, it is possible that the resolver has a different awareness of the name server capability during res_send processing than it had during res_init processing. The current awareness level is displayed before any queries are sent to the name servers.

In addition to the current EDNS0 capability of the name servers, if the resolver is monitoring name server responsiveness, the system-wide responsiveness statistics for each name server are also displayed. These statistics are based on a sliding 5-minute interval or a 30-second interval, so these numbers represent the volume of requests that are sent to the name server in the past 5 minutes or 30 seconds. Two statistics are maintained and displayed: the total number of requests that are sent to the name server by the resolver and the number of those requests that received no response from the name server. The resolver compares the percentage of failures for each name server against the value of the UNRESPONSIVETHRESHOLD resolver setup statement to determine whether a name server is to be considered unresponsive.

If the resolver is not monitoring name server responsiveness because UNRESPONSIVETHRESHOLD is set to 0, only the current EDNS0 capability information is displayed in the trace.

a If the AUTOQUIESCE operand is coded on the UNRESPONSIVETHRESHOLD statement, the resolver stops sending DNS queries that are generated by an application to unresponsive name servers. This message indicates whether resolver sends DNS queries to a particular name server. In these examples, all name servers are responding with a failure rate less than the UNRESPONSIVETHRESHOLD value, so the resolver sends DNS queries to all name servers, if necessary. If the name server was unresponsive, the resolver would generate a different trace message that indicates that the name server was to be skipped for DNS queries that are generated by an application.

b If the AUTOQUIESCE operand is not coded on the UNRESPONSIVETHRESHOLD statement, the statistics that are displayed represent 5 minutes of data, or the sum of five 1-minute statistical intervals. If the AUTOQUIESCE operand is coded, the statistics represent 30 seconds of data, or one 30-second interval of data. The number of intervals that are used to calculate the statistics is displayed on this trace entry.

22

Because the EDNS0 capability of the target name server is currently unknown, the resolver does not send the OPT RR record on the request to the name server. The DNS query that is built during res_mkquery processing is manipulated to remove the additional record.

a The receipt of the truncated UDP response causes resolver to reissue the query, to the same DNS, but this time the OPT RR record is not removed from the request data. This query is treated as an EDNS0 capability probe, and the response we receive from the DNS will determine whether future queries to the DNS include the OPT RR record or not. The transaction ID for this EDNS0 probe is different from the first query, and the new value is included in the trace information.

b As displayed on message line 21, the name server is considered to be up-level in terms of EDNS0 support, so the resolver sends the OPT RR record that res_mkquery built on the DNS query.

23

The receipt of a UDP package greater than 512 bytes, in response to an EDNS0 probe, indicates that this name server is up-level. Future communication with this name server will always include the OPT RR record to allow the name server to send UDP responses of greater than 512 bytes without requiring more EDNS0 probe requests.

24

Since the resolver included an OPT RR on the request to indicate that the resolver supports EDNS0, the name server also includes an OPT RR on the response. The DNS indicates it can accept up to 4096 bytes of data on a UDP message, but the resolver does not use this information.

25

The setting of Cache indicates that the system is currently performing resolver caching and the application is permitted to use the caching function. An individual application can be prevented from using the resolver caching function by specifying NOCACHE in the TCPIP.DATA data set.

26

Because resolver caching is active and available for use by this application, before the query is sent to the DNS, the resolver cache is queried to see whether any information is currently available for this resource.

a The cache has no information about the resource. The resolver must query the DNS to obtain the A record information.

b The cache has no AAAA record information about the resource, although it does have A record information (from the previous GetHostByName call). The resolver must query the DNS to obtain the AAAA information.

c The cache has A record information about the resource, which is saved from the previous GetHostByName call for the same resource, and the information is not expired, so the saved information is retrieved from the cache. No communication to the name server is necessary this time to obtain the A record information about the resource.

27

A timestamp is displayed when the query is actually sent to the DNS name server.

a A corresponding timestamp is displayed when the response from the DNS is successfully processed. If no response is received, the timestamp represents when the resolver stopped waiting for the response. The amount of time that elapses from the start of communication with the name server is also displayed.

28

Because resolver caching is active and available for use by this application, the information about this resource is cached for later reuse by this application or other applications on the system.

29

The resolver returns at most 35 IP addresses for a resource, even if, as in this case, there are more IP addresses that could be returned (20 IPv6 addresses, and 29 IPv4 addresses). Priority is given to IPv6 addresses, so in this case, all the possible IPv6 addresses are returned, but only 15 of the possible IPv4 addresses.

30

The time to live (TTL) value represents the amount of time that the resolver cache can use the returned information about the resource. The MAXTTL resolver setup statement can be used to define an upper limit on the actual TTL value that is used for a resource by the resolver.

31

The Type value indicates whether an IPv6 socket or an IPv4 socket is being used. IPv4 sockets are used when the system does not support IPv6 sockets or when only IPv4 addresses are specified for the name servers to be searched.

32

System name indicates the value of the system_name parameter and where it was obtained. If VMCF is running, the value is the name that VMCF used when it started. Otherwise, it is the value of the z/OS CVT (Communication Vector Table) CVTSNAME field.

33

At the completion of res_send processing, if UNRESPONSIVETHRESHOLD is non-zero, resolver updates the system-wide responsiveness statistics, and possibly the EDNS0 capability, for each name server that is contacted as part of res_send processing. In this example, resolver sent one request successfully to the first name server in the list, and as part of res_send processing discovered that the name server supported EDNS0. The EDNS0 probe that resolver sent, after the truncated response is received from the name server, is not included in the count of requests that are sent during res_send, since that probe is only used for resolver purposes.

Because only the first name server in the list was contacted, there were no updates to record for the second and third name servers in the list.

a Since resolver did not detect any changes in the EDNS0 capability of this name server, only the updates to the system-wide responsiveness statistics as part of res_send processing are displayed here.

34

The value of the UNRESPONSIVETHRESHOLD resolver setup statement at the time of the res_init processing is displayed. If the AUTOQUIESCE operand is specified on the resolver setup statement, the value that is displayed here would be in the form *percentage*,AUTOQUIESCE.

35

The resolver address space initializes even if syntax errors are detected in the

resolver setup file. This statement indicates whether errors are detected during resolver address space initialization and are not yet corrected.

36

This statement indicates whether the CTRACE TRACERES option is enabled. If the application enables the CTRACE TRACERES option, Trace Resolver output is collected as Resolver CTRACE records. For more information, see "CTRACE - RESOLVER."

37

The setting of CacheReorder indicates that the resolver reorders the list of cached IP addresses in a round-robin manner before the resolver returns the list to the requesting application. You can specify the NOCACHEREORDER statement in the TCPIP.DATA data set to disable cache reordering for an individual application.

Notes:

- If any errors occurred, see z/OS Communications Server: IP and SNA Codes.
- In a multitasking environment, if the LRECL of the Trace Resolver output is at least 128 characters, the MVS TCB storage address appears at the end of each line. The MVS TCB storage address can be useful in determining the origin of the resolver request.

CTRACE - RESOLVER

Component Trace (CTRACE) is used for the RESOLVER component (SYSTCPRE) to collect debug information. The TRACE RESOLVER traces information about a per-application basis and directs the output to a unique file for each application. The CTRACE shows resolver actions for all applications (although it might be filtered).

The CTRACE support allows for JOBNAME, ASID filtering, or both. The trace buffer is located in the Resolver private storage. The trace buffer minimum size is 128K, maximum 128M, default 16M. Trace records can optionally be written to an external writer.

The Resolver CTRACE initialization PARMLIB member can be specified at Resolver start time. Using the sample Resolver procedure shipped with the product, enter the following console command:

```
S RESOLVER,REUSASID=YES,PARMS='CTRACE(CTIRESxx)'
```

where *xx* is the suffix of the CTIRES*xx* PARMLIB member to be used. To customize the parameters used to initialize the trace, you can update the SYS1.PARMLIB member CTIRES00.

Note: In addition to specifying the trace options, you can also change the Resolver trace buffer size. The buffer size can be changed only at Resolver initialization.

If the CTIRES00 member is not found when starting the Resolver, the following message is issued:

```
IEEE538I CTIRES00 MEMBER NOT FOUND in SYS1.PARMLIB
```

When this occurs, the Resolver component trace is started with a buffer size of 16MB and the MINIMUM tracing option.

After Resolver initialization, you must use the TRACE CT command to change the component trace options (see Chapter 5, "TCP/IP services traces and IPCS support," on page 45). Each time a new component trace is initialized, all prior trace options are turned off and the new options are put into effect.

Trace options:

ALL
The standard set of CTRACE options for tracing resolver processing.

Rule: The TRACERES option is not included.

MINIMUM
The minimum set of options traces exceptions, Resolver initialization and termination, Resolver CTRACE changes, and Resolver operator messages.

TRACERES
Collects Trace Resolver information as Resolver CTRACE records.

The following sample shows the PARMLIB member.

```
/*********************************************************************/
/*                                                                   */
/*  IBM Communications Server for z/OS                               */
/*  SMP/E Distribution Name:  CTIRES00                               */
/*                                                                   */
/*  PART Name:  CTIRES00                                             */
/*                                                                   */
/*                                                                   */
/*  Copyright:                                                       */
/*               Licensed Materials - Property of IBM                */
/*               5694-A01                                            */
/*               (C) Copyright IBM Corp. 2001, 2003                  */
/*                                                                   */
/*                                                                   */
/*  Status:      CSV1R5                                              */
/*                                                                   */
/*                                                                   */
/*  DESCRIPTION = This parmlib member causes component trace for     */
/*               the TCP/IP provided Resolver to be initialized      */
/*               with a trace buffer size of 16M                     */
/*                                                                   */
/*               This parmlib member only lists those TRACEOPTS      */
/*               values specific to the TCP/IP Resolver. For a       */
/*               complete list of TRACEOPTS keywords and their       */
/*               values see:                                         */
/*               z/OS MVS INITIALIZATION AND TUNING REFERENCE.       */
/*                                                                   */
/*                                                                   */
/* $PARMS(CTIRES00),COMP(RES ),PROD(TCPIP ): Resolver Component Trace*/
/*                                          SYS1.PARMLIB member      */
/*                                                                   */
/*********************************************************************/
 TRACEOPTS
 /* ---------------------------------------------------------------- */
 /*   Optionally start external writer in this file (use both        */
 /*   WTRSTART and WTR with same wtr_procedure)                      */
 /* ---------------------------------------------------------------- */
 /*        WTRSTART(wtr_procedure)                                   */
 /* ---------------------------------------------------------------- */
 /*   ON OR OFF: PICK 1                                              */
 /* ---------------------------------------------------------------- */
           ON
 /*        OFF                                                       */
 /* ---------------------------------------------------------------- */
 /*   BUFSIZE: A VALUE IN RANGE 128K TO 128M                         */
```

```
/*            CTRACE buffers reside in Resolver Private storage    */
/*            which is in the regions address space.               */
/* ------------------------------------------------------------- *
         BUFSIZE(16M)
/*       JOBNAME(jobname1,...)                                     */
/*       ASID(Asid1,...)                                           */
/*       WTR(wtr_procedure)                                        */
/* ------------------------------------------------------------- */
/*   OPTIONS: NAMES OF FUNCTIONS TO BE TRACED, OR "ALL"            */
/* ------------------------------------------------------------- */
/*       OPTIONS(            */
/*               'ALL     '  */
/*               'TRACERES'  */
/*              ,'MINIMUM '  */
/*               )           */
```

When formatting the Resolver trace, use the CTRACE command. See Chapter 5, "TCP/IP services traces and IPCS support," on page 45 for the syntax for formatting a CTRACE. For the Resolver, the following formatting OPTIONS are available:

ASCII
Resolver trace data is displayed with ASCII translation only. The default is EBCDIC.

BOTH
Resolver trace data is displayed with both EBCDIC and ASCII translations. Each line of formatted data contains the offset, the hexadecimal display, the EBCDIC translation, then the ASCII translation. The default is EBCDIC.

EBCDIC
Resolver trace data is displayed with EBCDIC translation only. This is the default.

HEX
Resolver trace data is displayed only in hexadecimal (no ASCII or EBCDIC translation). The default is EBCDIC.

Guidelines:

- If the formatted CTRACE display wraps on the screen, use the IPCS PROFILE LINESIZE(*nn*) command, where *nn* is the largest number of characters that displays on one line.
- The choice of ASCII or EBCDIC translation does not apply to CTRACE TRACERES records.

CTRACE TRACERES

The CTRACE TRACERES option collects the same Trace Resolver output as described in "Interpreting the Trace Resolver output" on page 889, but the Trace Resolver records are formatted as Resolver CTRACE records. All TRACERES records have the same trace record name. The trace record name is Formatted Trace Resolver. TRACERES records might be interspersed among other Resolver CTRACE records when CTRACE ALL is also specified.

Guidelines:

- The CTRACE TRACERES and Trace Resolver functions can be used independently or together. When both functions are active, Trace Resolver

records are written as CTRACE records in addition to being written to the location that is described in "Specifying the Trace Resolver output location" on page 887.

- You must filter the recorded data to improve performance in the following ways:
  - Minimize the overhead of recording the trace
  - Make formatting faster
  - Save storage
  - Minimize wrapping, that is, new trace records overwriting older trace records

  Use the JOBNAME or ASID option when you capture the trace to limit the trace data to one application.

Enabling collection of CTRACE TRACERES records

Use one of the following methods to enable collection of Trace Resolver output as Resolver CTRACE records.

- Take the following steps to activate collection when you start the Resolver:
  1. Specify the CTRACE TRACERES option in the ctrace PARMLIB member.
  2. Start the Resolver with the PARMS keyword. For example:

     ```
     S resolver jobname, PARMS='ctrace parmlib member'
     ```
- Take the following steps to activate collection after the Resolver is started:
  1. Issue the **TRACE CT,ON** command and specify the CTRACE TRACERES option. For example:

     ```
     TRACE CT,ON,COMP=SYSTCPRE,SUB=(resolver jobname)
     R xx,OPTIONS=(TRACERES),END
     ```
  2. To collect CTRACE information for only a subset of applications, specify the JOBNAME, or ASID, or both options as filters. For example:

     ```
     TRACE CT,ON,COMP=SYSTCPRE,SUB=(resolver jobname)
     R xx,OPTIONS=(TRACERES),JOBNAME=(...),ASID=(...),END
     ```

Disabling collection of CTRACE TRACERES records

Issue the **TRACE,CT,ON,COMP=SYSTCPRE,SUB=(resolver jobname)** command and use one of the following responses to disable collection of Trace Resolver output as Resolver CTRACE records:

- If the Resolver CTRACE function was enabled without filters, respond with `R xx,OPTIONS=(),END` to restore the default CTRACE options.
- If the Resolver CTRACE function was enabled by using filtering, use one of the following responses to restore the default CTRACE options and remove the filter setting:
  - If you filter by using the JOBNAME option, respond with `R xx,OPTIONS=(),JOBNAME=(...),END`.
  - If you filter by using the ASID option, respond with `R xx,OPTIONS=(),ASID=(...),END`.
  - If you filter by using both JOBNAME and ASID options, respond with `R xx,OPTIONS=(),JOBNAME=(...),ASID=(...),END`.

  Guideline: In all cases, you can substitute OPTIONS=(MINIMUM) or OPTIONS=(ALL) for the OPTIONS=() response to disable the CTRACE TRACERES processing.

Formatting the CTRACE TRACERES option

Use the IPCS CTRACE,FULL command to format the trace regardless of whether the trace data is available in a dump or in an external data set. You can either type the CTRACE command on the IPCS command panel or use the panel interface. IPCS is also available in batch.

You can use several filters to limit the amount of formatted data. In addition to the CTRACE options such as start and stop time that IPCS provides, you can also use the following options:

JOBNAME
The job name of the application.

ASID The address space identifier of the application.

CID The connection identifier of the application with a specific res_init() function call. You can use CID to correlate related resolver calls from the same application.

Examples of the CTRACE TRACERES records

```
1  ================================================================00000031
2  MVSTST    TRACERES  000A0002  15:07:45.880242   Formatted Trace Resolver
3  ASID....     0024  TCB.... 007FF320  JOBN.... USER23    CID.... 00000222
4  Resolver Trace Initialization Complete -> 2014/01/07 15:07:45.861906

res_init Resolver values:
 Setup file warning messages = No
 CTRACE TRACERES option = Yes
 Global Tcp/Ip Dataset  = None
 Default Tcp/Ip Dataset = None
 Local Tcp/Ip Dataset   = //'SYS1.TCPPARMS(TCPDATA)'
 Translation Table      = CS390.STANDARD.TCPXLBIN
 UserId/JobName         = USER2
 Caller API             = LE C Sockets
 Caller Mode            = EBCDIC
 System Name            = MVSTST (from VMCF)
 UnresponsiveThreshold  = 25
 (L) DataSetPrefix = CS390
 (*) HostName      = MVSTST
 (L) TcpIpJobName  = TCPCS
 (L) Search        = tcp.raleigh.ibm.com
.
 (L) NsPortAddr    = 53            (L) ResolverTimeout    = 16
 (L) ResolveVia    = UDP           (L) ResolverUdpRetries = 1
 (*) Options NDots = 1
 (*) SockNoTestStor
 (*) AlwaysWto     = NO            (*) MessageCase        = MIXED
 (L) LookUp        = DNS
 (L) NoCache
res_init Succeeded
1  ================================================================00000037
2  MVSTST    TRACERES  000A0002  15:07:45.881099   Formatted Trace Resolver
3  ASID....     0024  TCB.... 007FF320  JOBN.... USER23    CID.... 00000222
4  res_init Started: 2014/01/07 15:07:45.881065
res_init Ended: 2014/01/07 15:07:45.881088
********************************************************************************
================================================================0000003E
```

The following list describes highlighted numbered areas in the CTRACE TRACERES records example:

1

This separator line shows the trace record number of the previous SYSTCPRE component in hexadecimal.

2

The first data line contains the host name (MVSTST), trace option (TRACERES), trace code (000A0002), time, and trace record name. The trace record name is Formatted Trace Resolver for all CTRACE TRACERES records.

3

The second data line contains the application's ASID (0024), the MVS TCB address (007FF320), the application's job name (USER23), and the connection identifier (00000222).

Tip: Use the ASID, JOBN, or CID value as input for the ASID, JOBNAME, or CID filter respectively to limit the amount of trace information that is formatted.

4

The contents of a CTRACE TRACERES record that follow the CTRACE headers represent a subset of the same Trace Resolver output as described in "Interpreting the Trace Resolver output" on page 889. The Trace Resolver output might be collected in multiple CTRACE records. In this example, two CTRACE records are necessary to collect all the Trace Resolver outputs that the res_init() processing generates.

Chapter 40. Diagnosing Simple Network Time Protocol (SNTP) problems

Simple Network Time Protocol (SNTP) is a standard protocol used to synchronize system clocks on routers and computer systems throughout the Internet through a specific formatted message. The Simple Network Time Protocol Daemon (SNTPD) is a TCP/IP daemon that is used to synchronize time between a client and a server.

This topic describes how to diagnose problems with SNTP daemon and contains the following sections:

- "Activating the SNTPD debug trace"
- "Abends"
- "Steps for stopping SNTPD"
- "Sample SNTPD debug output" on page 914

Activating the SNTPD debug trace

To activate the SNTPD debug trace, specify the -d or -df parameter when starting SNTPD via the z/OS UNIX shell or as an MVS started procedure.

If this option is used...	Then (phrase) . . .
-d parameter	Messages are written to the syslog daemon.
-df parameter **Restriction:** You must specify a path name and file name.	Messages are written to the file specified on the -df parameter.

Abends

An abend during SNTPD processing should result in messages and error-related information being sent to the system console. A dump of the error is needed unless the symptoms already match a known problem.

Steps for stopping SNTPD

If SNTPD was started from the z/OS UNIX shell, the **kill** command must be used to stop SNTPD.

Before you begin

Before you issue the **kill** command, you must determine the process ID (PID) of SNTPD.

Procedure

Perform the following steps to stop the PID of SNTPD:

1. To find the PID, use one of the following methods:
   - Use D OMVS,U=userid. (This is the USERID that started SNTPD from the shell.)
   - Use the **ps -ef** command from the shell.

© Copyright IBM Corp. 2000, 2015

- Write down the PID when you start SNTPD.

2. From a z/OS UNIX shell superuser ID, issue the **kill** command to the PID associated with SNTPD.

Results

You know you are finished when the following message appears: `EZZ9601I SNTP SERVER ENDED`. If SNTPD was started as an MVS started procedure, you must use the **stop** command to stop SNTPD. For example, code:

```
p sntpd
```

Sample SNTPD debug output

See z/OS Problem Management or see Chapter 3, "Diagnosing abends, loops, and hangs," on page 21, for information about debugging dumps produced during SNTP processing.

The following shows a sample of SNTP debug output.

```
Tue Apr  2 15:26:14 2002 SNTP enabled options: Opening debugging file /tmp/bc6.log
(Multicast: every 120 seconds) (PID FILE: /etc/sntpd.pid) (DEBUG FILE: /tmp/bc6.log
Tue Apr  2 15:26:14 2002 Writing PID to file /etc/sntpd.pid
Tue Apr  2 15:26:14 2002 EZZ9602I SNTP server initializing
Tue Apr  2 15:26:14 2002 Initializing signal handling
Tue Apr  2 15:26:14 2002 Set sigaction of signal SIGINT
Tue Apr  2 15:26:14 2002 Set sigaction of signal SIGTERM
Tue Apr  2 15:26:14 2002 Set sigaction of signal SIGABND
Tue Apr  2 15:26:14 2002 Set sigaction of signal SIGABRT
Tue Apr  2 15:26:14 2002 Set sigaction of signal SIGQUIT
Tue Apr  2 15:26:14 2002 Set sigaction of signal SIGHUP
Tue Apr  2 15:26:14 2002 Set sigaction of signal SIGTTOU
Tue Apr  2 15:26:14 2002 Initializing MVS command handling
Tue Apr  2 15:26:14 2002 Initializing pthread for MVS command
Tue Apr  2 15:26:14 2002 Initializing UDP socket(s)
Tue Apr  2 15:26:15 2002 SNTP port was set to 123
Tue Apr  2 15:26:15 2002 Bound to address: 9.67.2.1
Tue Apr  2 15:26:15 2002 Bound to address: 9.67.115.15
Tue Apr  2 15:26:15 2002 Bound to address: 9.67.2.2
Tue Apr  2 15:26:15 2002 Bound to address: 0.0.0.0
Tue Apr  2 15:26:15 2002 Initializing pthread for multicast/broadcast
Tue Apr  2 15:26:15 2002 Initializing pthread for unicast
Tue Apr  2 15:26:15 2002 EZZ9600I SNTP server ready
Tue Apr  2 15:28:15 2002 Sending NTP message to multicast address 224.0.1.1
Tue Apr  2 15:30:15 2002 Sending NTP message to multicast address 224.0.1.1
```

Chapter 41. Diagnosing Communications Server SMTP application problems

The Communications Server SMTP (CSSMTP) application transfers electronic mail from JES spool data sets to SMTP mail relays for delivery to the final destination. See z/OS Communications Server: IP Configuration Guide for overview and setup of CSSMTP.

The following information describes how to diagnose problems with the CSSMTP application.

- "Gathering diagnostic information"
- "Resolving initialization or logging problems" on page 916
- "Resolving SMTPNOTE CLIST problems" on page 917
- "Diagnosing and resolving Resolver problems" on page 918
- "Resolving problems from the JES spool data set" on page 919
- "Resolving mail problems" on page 923
- "Resolving MODIFY command problems" on page 933
- "Diagnosing checkpoint problems" on page 933
- "Monitoring resources used" on page 934
- "Monitoring target servers" on page 938

Gathering diagnostic information

You might need to collect multiple pieces of data to accurately diagnose problems.

- Capture the MVS console messages.
- If the problem is caused by configuration issues or target server definitions, then save the configuration file.
- Capture the CSSMTP application log messages. See "Steps for gathering log information" for details.
- If you get an abend during the CSSMTP application processing, messages and error-related information should be sent to the system console. A dump of the error is needed unless the symptoms already match a known problem. If an abend occurs, then save the resulting address space dump. See Chapter 3, "Diagnosing abends, loops, and hangs," on page 21 for details.

Steps for gathering log information

Gather log information.

Procedure

1. Determine where log data is located. If you are logging to a log file, then examine the LOGFILE DD statement in your started procedure to determine the log file or data set and save it. If you have not specified a LOGFILE DD, then the log will go to SYSLOGD.
2. If you are running multiple CSSMTP applications on the same system then you should be logging to SYSLOGD. Look in the SYSLOGD configuration file to determine where log records for CSSMTP are being written.
3. You should at least be executing with the default log level of error, warning, and event. If more logging information is needed, you can specify a loglevel

© Copyright IBM Corp. 2000, 2015

value greater than the default. See the LogLevel statement in z/OS Communications Server: IP Configuration Reference for valid CSSMTP / LogLevel statement information. The log level can also be changed with the MODIFY LOGLEVEL command, see the MODIFY command: Communications Server SMTP section in z/OS Communications Server: IP System Administrator's Commands for details.

4. If your LOGFILE DD is not a z/OS UNIX file system, you might observe buffering of the logfile output. Only regular memory files and UNIX file system files support the no buffering mode. A "//LOGFILE DD SYSOUT=*" will be treated as an MVS data set and may incur buffering. For further information, see z/OS XL C/C++ Runtime Library Reference, section General Description Controls buffering, for the specified stream, and z/OS XL C/C++ Programming Guide, section Buffering of C Streams.

Resolving initialization or logging problems

The following configuration messages indicate that initialization is done and you have successfully connected to at least one target for receiving mail.

- EZD1802I *csproc* INITIALIZATION COMPLETE FOR *extWrtName*
- EZD1821I *csproc* ABLE TO USE TARGET SERVER *ipAddress*

Table 95. Common initialization and logging problems

Symptoms	Problem/Cause	Action
BPXM023I	This message is prefixed to messages EZD1813E, EZD1815E, EZD1820E, and EZD1824E.	Define the CSSMTP id with READ access to BPX.CONSOLE CLASS(FACILITY), see z/OS UNIX System Services Programming: Assembler Callable Services Reference.
EZD1807I	This message indicates that initialization failed.	For details and other messages, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1824E	This message indicates that a TCP/IP stack is not available.	For details and other messages, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1815E	This message indicates that there was resolver problem at initialization.	For details and other messages, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1811I	This message indicates a problem in writing to the log file.	For details and other messages, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1835I	This message indicates that you did not configure a CHKPOINT DD statement.	For details and other messages, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).

Table 95. Common initialization and logging problems (continued)

Symptoms	Problem/Cause	Action
EZD1836I	This message indicates that you configured a CHKPOINT DD statement but checkpointing has failed.	For details and other messages, see z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).

Resolving SMTPNOTE CLIST problems

Table 96. Problems using SMTPNOTE CLIST to create mail messages on the JES spool data set

Symptoms	Problem/Cause	Action
EZA5579E	This message indicates that the mail messages generated from the SMTPNOTE CLIST is unable to be transmitted.	Ensure that the SMTPNOTE CLIST is customized. See z/OS Communications Server: IP Configuration Guide for details and see the appropriate EZA5579E message in z/OS Communications Server: IP Messages Volume 1 (EZA).
The mail message was queued to the JES spool queue, but it has not been processed by the CSSMTP application.	• The specified *smtpjob* in the SMTPNOTE CLIST may have the wrong name. • The specified *smtpjob* in the SMTPNOTE CLIST may not be active.	• Ensure that the SMTPNOTE CLIST is customized. See z/OS Communications Server: IP Configuration Guide for details. • Start the *smtpjob* application.

Diagnosing and resolving Resolver problems

Table 97. Diagnosing and resolving Resolver problems

Symptoms	Problem/Cause	Action
No IP addresses were resolved from the TargetServer configuration statements. EZD1815E (during initialization) is issued or EZD1845I (MODIFY REFRESH or MODIFY REFRESHIPLIST command) is issued. See the appropriate EZD1815E, EZD1845I messages in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).	Possible causes for this error are: • The Resolver is not initialized. • There may be a problem in the DNS or the Resolver setup. • The name specified on the TargetServer statement (TargetName or TargetMx parameter) may be incorrect.	• If the resolver is not initialized, start the resolver. See the appropriate EZD1815E, EZD1845I messages in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). • If the name on the TargetServer is in error, update the name and issue the MODIFY REFRESH command to reprocess the configuration file. • If the name on the TargetServer statement is correct, investigate your DNS and/or resolver setup.
Updated IP addresses were resolved from the TargetServer configuration statements with warnings. EZD1847I is issued. See the appropriate EZD1847I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).	Possible causes for this warning are: • A TargetName or TargetMx name could not resolve to any IP address • A TargetName or TargetMx name resolved to more than 4 IP addresses	Verify that the name specified on the TargetName or TargetMx parameter is correct. If the name is in error, update the name and issue the MODIFY REFRESH command to reprocess the configuration file. If the name is correct, investigate your DNS and/or resolver setup. If more than four addresses are resolved for a name specified on a TargetServer statement, only the first four are allowed and the rest ignored. **Note:** A warning is informational and no action is required.

Table 97. Diagnosing and resolving Resolver problems (continued)

Symptoms	Problem/Cause	Action
No changed IP addresses were resolved from the TargetServer configuration statements as a result of the MODIFY REFRESH command or MODIFY REFRESHIPLIST command. EZD1843I or EZD1844I is issued. See the appropriate messages in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).	If changed IP addresses were expected, then it might indicate a problem. • The name specified on the TargetServer statement may not be updated with the new name. • The name specified on the TargetServer statement was changed, but the updates in the DNS or the Resolver setup may not be refreshed.	Verify that the name specified on the TargetName or TargetMx parameter is correct. If the name is in error, update the name and issue the MODIFY REFRESH command to reprocess the configuration file. Refresh the resolver so that the response is not coming from the cache, then issue the MODIFY REFRESHIPLIST command to refresh the TargetName or TargetMx
MODIFY REFRESH or REFRESHIPLIST unsuccessful due to a MODIFY COMMAND in progress. EZD1806I is issued. See the appropriate EZD1806I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).	DNS setup may be in error causing the MODIFY COMMAND to take a long time.	Check if DNS is active at the IP address you specified for the NameServer or NSINTERADDR in the TCPIP.DATA data set. Check the port that DNS is listening on. The port is specified on the NSPORTADDR in the TCPIP.DATA data set.

The Resolver trace shows requests and responses sent to and received from name servers. It also shows if local hosts tables are used for name resolution. This trace helps you diagnose problems with host name resolution. See "TRACE RESOLVER" on page 884 for more information about how to activate Trace Resolver and collect output.

Resolving problems from the JES spool data set

Table 98. JES spool data set problems

Symptoms	Problem/Cause	Action
EZD1807I with JES NOT AVAILABLE	The CSSMTP application ends with the error during initialization because JES is not available. This is a permanent JES subsystem error during spool file initialization (for example, authentication errors, The CSSMTP application is not allowed to read it, IEFSSREQ failed).	For details and other messages, see the appropriate EZD1807I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).

Table 98. JES spool data set problems (continued)

Symptoms	Problem/Cause	Action
EZD1813I or EZD1816I A report may be generated to the MailAdministrator, if configured, or to the sysout file (see Figure 107 on page 922 for an example).	No mail messages in the spool file have been processed, because the entire JES spool file is permanently inaccessible. Possible causes for the error are: • There is a spool file allocation error • There is a spool file open error • This user ID is not permitted to send output to the external writer that is configured to the CSSMTP application. • The spool file size is larger than the JesJobSize.	For details and other messages, see the appropriate EZD1813I or EZD1816I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). For details on MailAdministrator and JesJobSize configuration options, see z/OS Communications Server: IP Configuration Reference. For details on JES issues relating to CSSMTP see z/OS Communications Server: IP Configuration Guide.
EZD1813I or EZD1816I A report may be generated to the MailAdministrator, if configured, or to the sysout file (see Figure 107 on page 922 for an example).	Some of the mail messages in the JES spool file have been processed and sent, but the rest of the JES spool file is permanently inaccessible. Possible causes for the error are: • IO error during read • Permanent SAPI errors during read • Syntax errors from JES spool processing, for example: – No valid HELO or EHLO SMTP command – There is an incorrect beginning SMTP command for HELO, EHLO, MAIL FROM – There is an incorrect ending SMTP command ".", RSET or next HELO or EHLO – There are more than *JESSyntaxErrLimit* invalid SMTP commands, or sequence errors – The CSSMTP application security user exit fails with the return code to stop processing the JES file	For details and other messages, see the appropriate EZD1813I or EZD1816I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). For details on MailAdministrator options and user exits see z/OS Communications Server: IP Configuration Reference. For details on JES issues relating to CSSMTP see z/OS Communications Server: IP Configuration Guide. For details on syntax of mail see z/OS Communications Server: IP User's Guide and Commands.

Table 98. JES spool data set problems (continued)

Symptoms	Problem/Cause	Action
EZD1813I or EZD1816I A report may be generated to the MailAdministrator, if configured, or to the sysout file (see Figure 107 on page 922 for an example).	Some of the individual mail message in a JES spool file are permanently non-deliverable. Possible causes for errors are: • There are non-mail boundary errors, for example: – Missing RCPT TO – JesMsgSize error • The CSSMTP application security exit fails with the return code that causes only individual mail message to fail.	For details and other messages, see the appropriate EZD1813I or EZD1816I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). For details on MailAdministrator options and user exits see z/OS Communications Server: IP Configuration Reference. For details on JES issues relating to CSSMTP see z/OS Communications Server: IP Configuration Guide. For details on syntax of mail see z/OS Communications Server: IP User's Guide and Commands.

The following is an error report example to MailAdministrator or to the sysout file.

```
[1]                      Error Report for USER1U (JOB00085)

[2]  Job USER1U  /        /DEST     (JOB00085) created by VIC000.USER1 at Thu, 21 Aug 2008 12:41:24 -0400
[2]  For DDname: SYSUT2    Dataset name: USER1.USER1U.JOB00085.D0000102.?
[3]  CSSMTP_XYZ generated the following messages:

[4]  --- Line    14 Mail     2 : Undeliverable mail for testid@test.com
       Message-Id:  USER1U.JOB00085.VIC000@tcp.com.Aug212008.124124.808675.2>
       Error  : No target server capable of receiving mail:
   mail extensions not supported.

       Mail was not delivered to the following recipients:
       user41@vic000.tcp.com
[5]  --- Line    54 Mail     5 : RCPT TO:   <user@vic000.tcp.com>
   501 5.1.1 JES Syntax error in mailbox '<user'
[6]  --- Line    40 Mail     4 : Undeliverable mail for userx@vic000.tcp.com
       Message-Id:<USER1U.JOB00085.VIC000@tcp.com.Aug212008.124124.808675.4>

       Mail was not delivered to the following recipients:
       usery@vic000.tcp.com
       Reply  : 550 User 'usery' Unknown
[7]  --- Line    26 Mail     3 : Undeliverable mail for testid@test.com
       Message-Id:<USER1U.JOB00085.VIC000@tcp.com.Aug212008.124124.808675.3>

       Mail was not delivered to the following recipients:
       userx@vic000.tcp.com
       Reply  : 550 User 'userx' Unknown

[8]  --- Line    48 Mail     6 : Undeliverable mail for testid@example.com
       Message-Id:<USER1U.JOB00085.VIC000@example.com.Aug212008.124124.808675.6>
       Error  : No target server capable of receiving mail:
    unable to establish a TLS connection.
       Mail was added to extended retry

       Mail was not delivered to the following recipients:
       user2@vic000.example.com
[9]  Completed at Thu, 21 Aug 2008 12:41:33 -0400

[10]  9 = mail messages found
[11]  1 = mail messages with errors

[12]  3 = recipients to whom mail was sent successfully
[13]  4 = recipients to whom mail messages could not be delivered


[14]  Disposition of the JES file was HOLD
```

Figure 107. Error report example

Notes:

1. This identifies the report with the jobname and jobid that contained the JES spool file
2. These further identify the source of the JES spool file with the jobname, the procedure step name, the job step name, the JES jobid, the originator node user ID, and the local time zone.
3. The CSSMTP application jobname and external writer name
4. The line and mail found in the JES spool file cannot be delivered to listed recipients below. The Message-ID is from the header line which identifies this mail to the other SMTP servers. The Error: lines describe the reason the CSSMTP application could not send the mail.

5. The line and mail found in the JES spool file where a JES syntax error was found. The next line identifies the cause of the error. In this case a blank in the mailbox.
6. The line and mail found in the JES spool file cannot be delivered to listed recipients below. In this case the target server rejected the mail because a mailbox could not be accepted.
7. Notice that the mail is listed in the order that they were completed and not in the order that they appear in the spool file.
8. The line and mail found in the JES spool file cannot be delivered. The mail message was saved for extended retry.
9. The processing of the spool file was completed at this date and time.
10. The number of mail messages found in the spool file.
11. The number of mail messages that contained parsing errors.
12. The number of recipients to whom mail sent successfully.
13. The number of recipients to whom mail could not be delivered.
14. The disposition of the JES spool file that was processed. If the value is HOLD or DELETE it reflects the setting of the BadSpoolDisp configuration statement. If the value is KEEP, it indicates that the CSSMTP application was interrupted during processing of this file. The file will continue to be processed when the application is restarted.

Resolving mail problems

To resolve common problems with connecting to or communicating with a target server at a certain IP address, begin by reviewing messages sent to the system console. Review the CSSMTP log file for more detail about the specific problem. Table 99 shows the possible messages and their meaning.

Target server problems

Table 99. Common messages indicating target server status

Symptoms	Problem/Cause	Action
EZD1817I	This message indicates that the application has been unable to successfully connect to the target server at the indicated IP address due to a socket connect() failure or timeout.	For details and other messages, see the appropriate EZD1817I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).

Table 99. Common messages indicating target server status (continued)

Symptoms	Problem/Cause	Action
EZD1818I	This message indicates that the application has been unable to communicate successfully with the target server at the indicated IP address due to an SMTP protocol problem or timeout after a successful connect. SMTP protocol problems can be the following: • Initial SMTP greeting not sent • 4xx or 5xx reply codes sent on EHLO SMTP command • 4xx or 5xx reply codes sent on HELO SMTP command Other problems are socket read() or write() failures on the connection.	For details and other messages, see the appropriate EZD1818I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). If the problem appears to be in communicating with a target server, setting the loglevel to at least 39 will help to capture TCP/IP traces that show SMTP commands and remote SMTP server replies between the CSSMTP application and the TCP/IP network.
EZD1819I	This message indicates that the application has been unable to establish a TLS connection with the target server at the indicated IP address due to an SMTP protocol problem, a security certificate problem or a timeout. SMTP protocol problems can have the following causes: • STARTTLS option not supported on EHLO SMTP command • 4xx or 5xx reply codes sent on STARTTLS SMTP command • The stack does not support TLS policies. • The policy agent is not started resulting in SIOCTTLSCTL ioctl() failures. • Socket read() or write() failures on the connection.	For details and other messages, see the appropriate EZD1819I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). For details setting up TLS security on the CSSMTP application, see z/OS Communications Server: IP Configuration Guide. If TLS is enabled and the definitions are configured on the client stacks but EZD1819I and log file ERROR messages indicates this is due to security see Chapter 28, “Diagnosing Application Transparent Transport Layer Security (AT-TLS),” on page 707. If TLS is set up correctly on CSSMTP, then verify that TLS is set up correctly on the server.

Table 99. Common messages indicating target server status (continued)

Symptoms	Problem/Cause	Action
EZD1820E	This message indicates that no target server is capable of receiving mail. There may be a networking problem affecting all target servers. Processing of mail is suspended when this message is generated.	For details and other messages, see the appropriate EZD1820E message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1824E	This message indicates that the application has been unable to successfully communicate to the TCP/IP stack as the result of a socket() failure. Processing of mail is suspended when this message is generated.	For details and other messages, see the appropriate EZD1824E message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).

Mail problems

Generally, common problems dealing with individual mail messages are handled by creating an undeliverable mail notification when the mail message becomes undeliverable.

- If ReturnToMailFrom is set to YES on the UNDELIVERABLE statement or the default value is used, an undeliverable mail notification is created that contains the original mail text as well as additional information that indicates the reason why the mail could not be delivered. The undeliverable mail notification is sent to the originator's mail address as specified on the MAIL FROM command.
- If requested, a report is generated that contains the error text indicating why the mail message could not be delivered. The report is handled according to the setting on the REPORT configuration statement. See z/OS Communications Server: IP Configuration Reference for details.

The following steps help in debugging mail problems.

1. Start by reviewing the error text captured in the undeliverable mail notification for the undeliverable mail message. The following list contains reasons why mail can become undeliverable.
   - Common problems with mail messages can be that there are no target servers with the correct capabilities to send the mail, such as the following:
     - The mail requires a TLS connection. This requires the target server to be an ESMTP which supports the EHLO SMTP command and the extension option STARTTLS. The target server must reply positively to the STARTTLS command (2xx reply code). Also the CSSMTP application must be able to successfully issue the socket ioctl() to communicate with the TCP/IP stack to establish the TLS connection.
     - The mail message size is larger than what the target server indicates it can support on the EHLO reply response when it is an ESMTP server.
     - The mail message size is larger than what is configured on the parameter MessageSize associated with the TargerServer statement when it is an SMTP server.

- Target server replies to SMTP commands with a 4xx reply code, which indicates that the mail message send be retried. However, the number of retries exceeded the configured maximum, see the statement RetryLimit in z/OS Communications Server: IP Configuration Reference for more information.
- Target server connection timeouts occur, which indicates that the mail message send be retried. However, the number of retries exceeded the configured maximum, see the statement RetryLimit in z/OS Communications Server: IP Configuration Reference for more information.
- Target server replies to MAIL FROM, RCPT TO, DATA and sending end of mail (EOM) sequence with a 5xx reply code indicating failure.

The following example shows general error text for a single recipient in the undeliverable mail notification.

```
[1]  Received: from host1.ibm.com (host1)
  by host2.ibm.com (host2 [x.x.x.x])
  for <USER1@host1>
  with ESMTP (IBM CSSMTP z/OS V01R11.00)
  Id SMTP2.JOB00051.host2@ibm.com.Nov192008.134147.442406.1U ;
 Wed, 19 Nov 2008 13:54:31 -0500
[2]  Date: Wed, 19 Nov 2008 13:54:31 -0500
[3]  From: <CSSMTP@host2.ibm.com>
[4]  To: <user1@ibm.com>
[5]  Subject: Undeliverable Mail for MSG-ID SMTP2.JOB00051.host2@ibm.com.Nov192008.134147.442406.1

[6]  Error text: No target server capable of receiving mail:
retry limit count exceeded for this mail. Tried (1).

[7]  Mail was not delivered to the following recipients:
[8]    <user2@samehost.ibm.com>
[9]    <user3@samehost.ibm.com>
[10]   Reply text:550 Recipient unknown
[11]   <user4@samehost.ibm.com>
[12]   <user5@badhost.ibm.com>
[13]   Reply text:550 "badhost" not found



[14]  Original mail text:
   Received: from host1.ibm.com (host1)
     by host2.ibm.com (host2 [x.x.x.x])
     for <USER1@host1>
     with ESMTP (IBM CSSMTP z/OS V01R11.00)
     Id SMTP2.JOB00051.host2.ibm.com.Nov192008.134147.442406.1 ;
    Wed, 19 Nov 2008 13:53:18 -0500
   FROM:     user1@ibm.com
   TO:       user2@samehost.ibm.com, user3@samehost.ibm.com, user4@samehost.ibm.com, user5@badhost.ibm.com
   SUBJECT: MSG1 JOB SMTP2 EHLO
   Date: Wed, 19 Nov 2008 13:53:18 -0500
   Message-ID: <SMTP2.JOB00051.host2@ibm.com.Nov192008.134147.442406.1>

   This is only a test.
```

Figure 108. Example of undeliverable mail notification

[1]
Receive line for the undeliverable mail notification.

[2]
Date and time the undeliverable mail notification was created.

[3]
The email address of the CSSMTP application that created the undeliverable mail notification.

[4]
The email address to whom the undeliverable mail notification is sent.

[5]
The subject line showing the MSG-ID SMTP2.JOB00051.host2@ibm.com.Nov192008.134147.442406.1 for the original piece of mail. This is the same value as the Message-ID header in the original mail text.

[6]
General error text applies to all recipients that mail was not delivered to that do not have a specific error reply text. In this example, the general error text applies to the recipients in line [8] and [11].

[7]
After this line, the list of recipients that did NOT receive the original mail text.

[9]
Original mail text was not delivered to this recipient <user3@samehost.ibm.com> because of the specific error described in line [10] reply text. This recipient was unknown to the target server.

[12]
Original mail text was not delivered to this recipient <user5@badhost.ibm.com> because of the specific error described in line [13] reply text. The email address of the recipient was not acceptable to the target server.

[14]
The original mail message - headers and content.

Table 100. Common mail text problems

Symptoms	Problem/Cause	Action
No target server capable of receiving mail: original message size too large.	The mail size is too large. Check size of the mail including headers. Use the MODIFY DISPLAY TARGETS command, see z/OS Communications Server: IP System Administrator's Commands to know what mail size is supported by target servers and whether the target server is SMTP or ESMTP.	If appropriate reduce mail size and resend. Otherwise, increase the mail size supported by the target server. To check target server configuration, see TargetServer statement in z/OS Communications Server: IP Configuration Reference if SMTP. Correct the target server on remote host if it is ESMTP.
No target server capable of receiving mail: retry count or interval is zero for this mail.	Review the log file. The target server is using reply code 4xx or is timing out the connection which requires a retry to occur.	Correct the target server or add a new target to the configuration. Change statement RetryLimit to allow retries to occur.

Table 100. Common mail text problems (continued)

Symptoms	Problem/Cause	Action
No target server capable of receiving mail: retry limit count exceeded for this mail.	Review the log file. The target server is using reply code 4xx or is timing out the connection which requires a retry to occur.	Correct the target server or add a new target to the configuration.
No target server capable of receiving mail: TLS support not available.	The batch job containing this mail required a TLS connection. The STARTTLS command is part of the batch job. Use the MODIFY DISPLAY TARGETS command. See z/OS Communications Server: IP System Administrator's Commands to determine whether TLS is supported by the target servers. Target servers must be ESMTP to support TLS.	If appropriate remove STARTTLS command from batch job. Check the target server configuration. Review the log file for the reply code from the target server or connection timeout which requires a retry to occur. Correct the target server or add a new target server capable of doing TLS to configuration.

Table 101. Dead letters problems

Symptoms	Problem/Cause	Action
EZD1826I	This message indicates that the dead letters cannot be written to the configured or default dead letter directory.	For details and other messages, see the appropriate EZD1826I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). No more dead letters are stored. The CSSMTP application internally sets the DeadLetterAction statement as it was set to Delete. The user can automate on this message to monitor storage utilization, and decide whether MODIFY SUSPEND command should be issued to suspend all new spool file processing and should examine and clean up all stored and unneeded dead letters to free up storage. Then MODIFY REFRESH command should be issued with DeadLetterAction set to STORE to tell the CSSMTP application to write dead letters into the configured or default dead letter directory again. MODIFY RESUME command is issued (if MODIFY SUSPEND command was issued previously) to resume new spool file processing.

Table 102. Mail Directory problems

Symptoms	Problem/Cause	Action
EZD1827I	This message indicates that mail messages cannot be written to the configured or default mail directory associated with the ExtendedRetry statement.	For details and other messages, see the appropriate EZD1827I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). No more extended retry mail messages are stored. The CSSMTP application ends. Look up the return code and reason code values associated with the console message to determine the next action to take. If you determine that the z/OS UNIX file system mail directory is not accessible in write mode, perform one of the following actions: • Mount the z/OS UNIX file system in write mode. • Specify a new directory on the MailDirectory parameter of the ExtendedRetry configuration statement in the CSSMTP configuration file. CSSMTP must have write access to this directory. Restart the CSSMTP application.

Table 102. Mail Directory problems (continued)

Symptoms	Problem/Cause	Action
EZD1852I	This message indicates that mail messages cannot be stored to the configured or default mail directory associated with the ExtendedRetry statement because the file system is out of space.	For details and other messages, see the appropriate EZD1852I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). No more extended retry mail messages are stored. The CSSMTP application ends. Perform one of the following actions and restart the CSSMTP application: • Delete any files in the file system that are no longer needed. • Define a new mail directory on z/OS UNIX file system with more space: 1. Define a new larger file system. 2. Mount the file system. 3. Use the **mkdir** command to define a new mail directory. 4. Copy all the files from the old mail directory to the new mail directory. 5. Update the configuration to use the new mail directory.

Using the undeliverable mail notification for problem determination

- If the error text captured in the undeliverable mail notification for the undeliverable mail is not enough to explain the problem, then review the log file that is generated under the started task by using the message-ID value that is found in the Subject: header and search through the log file to see how the mail message was processed.
- If the problem appears in communicating with a target server, set the loglevel value to at least 39 to capture TCP/IP traces. This shows SMTP commands and remote SMTP server replies between the CSSMTP application and the IP network. Resend the mail message.
- If the mail does not get sent and no undeliverable mail notification is returned to the originator of the mail, wait for CSSMTP to retry to send the mail message based on the configured RetryLimit statement values.

  If there is still no undeliverable mail notification, do the following checks:

- Check the configuration file for the undeliverable statement DeadLetterAction. If the statement is set to DELETE, then the undeliverable mail notification was deleted. If DeadLetterAction is set to STORE, then check the z/OS UNIX file system under the configured or default directory path.
- Check the log file that was generated under the started task, by using the undeliverable mail notification message ID value to determine how the mail was processed. The following is the message ID for the undeliverable mail notification:

  ```
  SMTP2.JOB00051.host2@ibm.com.Nov192008.134147.442406.1U
  ```

Duplicate mail messages

If you run CSSMTP with SMTPD to test compatibility of your mail workload, ensure that you have set TestMode to YES in the CSSMTP configuration file. This parameter cannot be changed dynamically. If you do not set the parameter, CSSMTP delivers mail messages that do not contain errors, resulting in duplicate mail messages being delivered if SMTPD also runs and delivers mail messages.

Verifying mail workload compatibility with SMTPD

SMTPD will be removed from z/OS and CSSMTP in the future. You can use the TestMode feature of CSSMTP with EZBMCOPY to verify whether your mail workload is compatible with CSSMTP before you make the switch. For more information, see Withdrawal of SMTPD and sendmail in z/OS Communications Server: IP Configuration Guide.

When you run CSSMTP with TestMode set to YES, the only output that CSSMTP produces is error messages for noncompliant mail messages. Because CSSMTP uses its normal error reporting mechanisms to report these errors, ensure that you configure a valid destination for CSSMTP error reports. You must have a valid destination for these reports and search the reports in the correct place to ensure that you can detect any mail workload that is not compatible with CSSMTP.

When CSSMTP running with TestMode set to YES completes a JES spool file, CSSMTP deallocates the file if no errors are found. If errors are found, it handles the spool files according to the setting of the BADSPOOLDISP configuration statement.

If you run CSSMTP with SMTPD to test compatibility of your mail workload, ensure that you set TestMode to Yes in the CSSMTP configuration file. This parameter cannot be changed dynamically, so you must stop and restart CSSMTP if it is set incorrectly. If you do not set the parameter, CSSMTP delivers mail messages that do not contain errors, resulting in duplicate mail messages being delivered if SMTPD also runs and delivers mail messages.

Resolving MODIFY command problems

Table 103. Common MODIFY command

Symptoms	Problem/Cause	Action
EZD1806I	This message indicates the modify command was unsuccessful.	For details and other messages, see the appropriate EZD1806I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1839I	This message indicates MODIFY REFRESH completed with errors.	For details and other messages, see the appropriate EZD1839I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1845I	This message indicates that there was a resolver problem with the MODIFY REFRESH or MODIFY REFRESHIPLIST commands.	For details and other messages, see the appropriate EZD1845I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).

Diagnosing checkpoint problems

Since each message can have several specific problems, reviewing the trace logs captured under the started task will provide more details on common problems dealing with the checkpoint functions. The following table shows the possible messages and their meaning.

Table 104. Common messages indicating checkpoint status

Symptoms	Problem/Cause	Action
EZD1835I	This message indicates that the CHKPOINT DD statement was missing or the data set name was 'NULLFILE'. The function is disabled.	For details and other messages, see the appropriate EZD1835I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
EZD1836I	This message indicates that the checkpoint data set could not be opened.	For details and other messages, see the appropriate EZD1836I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).

Table 104. Common messages indicating checkpoint status (continued)

Symptoms	Problem/Cause	Action
EZD1849I	This message indicates that JES spool file could not be allocated by the SAPI interface. The spool file may have been previously purged by an operator command.	For details and other messages, see the appropriate EZD1849I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD).
Abend S08B	This abend is issued by the MVS Data-In-Virtual component when accessing the checkpoint data set. Reason code xxxx001D. The checkpoint data set is not a valid VSAM linear data set or an incorrect control interval size was used when the VSAM data set was defined.	For details about the S08B abend, see z/OS MVS System Codes.

Monitoring resources used

After the CSSMTP application is started it will periodically check the following resources, and determines if any action needs to be taken:

- JES tasks
- Storage use in the CSSMTP application address space
- Storage use for z/OS UNIX file system in the configured or default dead letter directory
- Storage use for z/OS UNIX file system in the configured or default mail directory

Table 105. Common messages used for monitoring resources

Symptoms	Problem/Cause	Action
EZD1856I	This message indicates either that over 75% of DEST JES tasks and 75% of WRITER JES tasks are waiting for long retry processing to complete one or more mail messages in a spool file.	For details and other messages, see the appropriate EZD1856I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). You should automate on this message to monitor the JES tasks that are not available. Use automation to decide whether to reissue the MODIFY FLUSHRetry command to remove mail messages from the retry queue to free up the JES tasks, or to issue the MODIFY SUSPEND command to suspend all new spool file processing. Message EZD1857I is issued when the usage of DEST JES tasks and the usage of WRITER JES tasks that are waiting for long retry processing drops below 50%. You can automate this message to monitor the usage of JES tasks, and to decide whether you should issue the MODIFY RESUME command (if MODIFY SUSPEND command was issued previously) to resume processing for the new spool files.

Table 105. Common messages used for monitoring resources (continued)

Symptoms	Problem/Cause	Action
EZD1858I	This message indicates that storage utilization in the CSSMTP application address space exceeds 75%.	For details and other messages, see the appropriate EZD1858I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). You should automate on this message to monitor storage utilization. Use automation to decide whether to issue the MODIFY SUSPEND command to suspend all new spool file processing. EZD1859I is issued when storage use drops below 50% in the CSSMTP application address space. You can automate this message to monitor storage use, and to decide whether you should issue the MODIFY RESUME command (if MODIFY SUSPEND command was issued previously) to resume processing for the new spool files.

Table 105. Common messages used for monitoring resources (continued)

Symptoms	Problem/Cause	Action
EZD1860I	This message indicates that storage utilization for z/OS UNIX file system in the configured or default dead letter directory exceeds 75%.	For details and other messages, see the appropriate EZD1860I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). You should automate on this message to monitor storage use. Use automation to decide whether to issue the MODIFY SUSPEND command to suspend all new spool file processing and clean up all stored and unneeded dead letters to free up storage. EZD1861I is issued when storage use for the z/OS UNIX file system in the configured or default dead letter directory drops below 50%. You can automate this message to monitor storage use, and to decide whether you should issue the MODIFY RESUME command (if MODIFY SUSPEND command was issued previously) to resume the processing for the new spool files.

Table 105. Common messages used for monitoring resources (continued)

Symptoms	Problem/Cause	Action
EZD1862I	This message indicates that storage utilization for z/OS UNIX file system in the configured or default mail directory associated with the ExtendedRetry statement exceeds 75%.	For details and other messages, see the appropriate EZD1862I message in z/OS Communications Server: IP Messages Volume 2 (EZB, EZD). You should automate on this message to monitor storage use. Use automation to delete any files in the file system that are no longer needed. Do not delete any files from the mail directory. If target servers are available to receive the mail, decide whether to issue the MODIFY FLUSHRtry,AGE=*days* command to delete old mail messages from the mail directory. Use the MODIFY DISPLAY,TARgets command to monitor the number of mail messages in the extended retry directory and the state of the target servers used by CSSMTP. EZD1863I is issued when storage use for the z/OS UNIX file system in the configured or default mail directory drops below 50%. You can automate on this message to monitor the storage use.

Monitoring target servers

The Communication Server SMTP (CSSMTP) application automatically monitors the target servers for certain failures when the application is unable to communicate with the servers. Console messages (messages EZD1817I, EZD1818I, and EZD1819I) are generated for each target server. In some cases, the target servers can communicate with the CSSMTP application but the target servers cannot accept the mail message due to temporary error conditions. For temporary error conditions, the target servers pass the reply code 4*xx*. See RFC 2821 for more information about reply code processing. See Appendix D, "Related protocol specifications," on page 1003 for information about how to access RFCs.

To understand and handle the temporary error conditions, perform the following steps:

1. Determine how the CSSMTP application handles temporary error conditions.

Two configuration statements control how CSSMTP handles temporary error conditions:

RetryLimit
Mail messages are saved in memory and wait to be retried. The JES spool data set is held and waits for mail messages to be sent or failed permanently. This state is called long retry.

ExtendedRetry
Mail messages are saved in the z/OS UNIX file system and wait to be retried. The JES spool data is deleted if there are no other errors to report in the spool data set. This state is called extended retry.

Use the MODIFY DISPLAY,CONFIG command to determine the current value of these statements.

2. Monitor whether temporary errors are occurring. Use the MODIFY DISPLAY,TARgets command to determine the number of mail messages currently in retry processing and check the following cases:
   - Check whether the number is increasing or is high under the GLOBAL INFORMATION: CURRENT RETRY label for mail messages in long retry.
   - Check whether the number is increasing or is high under the EXTENDED RETRY: CURRENT label for mail messages in extended retry.
3. Use traces to determine the temporary errors. The target servers must be available.
   a. Issue the MODIFY LOGlevel,LEVEL=logLevel command. Set the logLevel value to the current level plus 32, which is the value at which message-level TCP/IP messages are logged. This command traces the CSSMTP commands and remote SMTP target server replies between CSSMTP and the TCP/IP network.
   b. Look for the remote SMTP server replies that contain any of the following information:
      - A reply code that begins with of the number 4, for example, `450`.
      - A reply message that indicates a timeout with the servers.

   If mail messages are not being sent when traces are activated, you can force the mail messages off the retry queues and return to the active queue.
   - For mail messages in long retry, use the MODIFY FLUSHRetry,TKID=*tkid* command, where *tkid* is the task ID associated with the JES spool data set.
   - For mail messages in extended retry, use the MODIFY FLUSHRetry,AGE=*days* command.
4. Handle the temporary error conditions.
   - If the target server is short of resources such as memory or space, take one of the following actions:
     – Add resources to existing target servers.
     – Add a target server that has extra resources to the target server list in CSSMTP configuration file. The maximum number of possible target servers is four. Use the MODIFY REFRESH command to make the CSSMTP application update the target server list.
   - If the recipient's mailbox is full, take corrective actions at the target server where the mailbox is located.
   - If timeout errors occur during the communication between CSSMTP and target servers, check the configured timeout values in the CSSMTP configuration file and the configured timeout values at the target servers, and increase the values.

ExtendedRetry Mail Directory

The CSSMTP application has an extended retry function. When this function is activated, it creates a control file (cf) and a data file (df) in a directory that is specified by the MailDirectory parameter that is associated with the ExtendedRetry statement. Users or mail administrators cannot modify the content of these files because the CSSMTP application needs to control the content of these files and to control when they are created and deleted. If the CSSMTP application has a problem with these files, it renames them to CF and DF. You can check the CSSMTP log to determine why the CF and DF files were renamed. You can delete them after you complete your analysis.

Chapter 42. Diagnosing storage abends and storage growth

The key to the successful resolution of most storage problems is to first determine whether the storage problem you are experiencing is related to common, private or communication storage manager (CSM) storage. This topic outlines steps you can use to determine the type of storage you are having a problem with and the steps to take to diagnose the storage problem.

This topic contains the following subtopics:

- "Storage definitions"
- "Monitoring storage utilization" on page 942
- "Limiting TCP/IP common and private storage utilization" on page 943
- "Limiting CSM storage utilization" on page 943
- "Storage messages" on page 944
- "Abends" on page 945
- Problem determination
- Collecting documentation

Storage definitions

TCP/IP uses several types of storage.

common storage
Common storage is shared across the whole system and can be accessed from any address space. Common storage is managed by the Virtual Storage Management component of the z/OS operating system. TCP/IP's usage of common storage for the most part is for ECSA (extended common service area).

private storage
Private storage is storage that is unique to an address space. Private storage is also referred to as pool storage.

CSM Communication Storage Manager (CSM) enables TCP/IP, VTAM, and other applications to use CSM buffers to reduce data movement. CSM interfaces with the z/OS operating system to provide the storage buffers. Buffers are maintained in both common storage and in the CSM data spaces. The application (for example TCP/IP) has the option of requesting the buffers from common storage or the CSM data spaces. The storage is managed in pools of the following predefined buffer sizes:

- 4 KB
- 16 KB
- 32 KB
- 60 KB
- 180 KB

See z/OS Communications Server: CSM Guide for additional information relating to the CSM component.

© Copyright IBM Corp. 2000, 2015

Monitoring storage utilization

Storage is a resource that many users monitor very closely to determine their average utilization. You can monitor storage by using storage monitors; by manually issuing TCPIP and CSM display commands; or by using NetView CLISTs that are triggered at specific time intervals. Automatically issuing commands to your system log at periodic intervals is the most efficient way to monitor your storage utilization.

Use this log output to establish a storage utilization history. Knowing how much common, private, or CSM storage you typically use can be helpful when trying to resolve problems where TCP/IP storage utilization is increasing (also referred to as a storage creep) or you receive an abend related to an out of storage condition.

Monitor TCP/IP common and private storage utilization by issuing the Display TCPIP,,STOR console command:

For example:

```
d tcpip,,stor
TCPIP STORAGE
TCPCS    STORAGE  CURRENT  MAXIMUM    LIMIT
TCPCS    ECSA         14M      28M     120M
TCPCS    POOL         52M      62M  NOLIMIT
DISPLAY TCPIP STOR COMPLETED SUCCESSFULLY

or

d tcpip,tcpip2,stor
TCPIP STORAGE
TCPIP2   STORAGE  CURRENT  MAXIMUM    LIMIT
TCPIP2   ECSA      45654K   56823K  204800K
TCPIP2   POOL     124634K  143743K  524288K
DISPLAY TCPIP STOR COMPLETED SUCCESSFULLY
```

For additional information regarding the Display TCPIP,,STOR command, see z/OS Communications Server: IP System Administrator's Commands.

TCP/IP's CSM storage utilization is not included in the TCP/IP display. TCP/IP's CSM storage utilization can be monitored by issuing the following console commands.

To display information about storage managed and used by CSM for all owners, issue the following command:

d net,csm,ownerid=all

To display information about CSM utilization for TCP/IP:

d net,csm,ownerid=TCPIP asid

For example:

```
d net,csm,ownerid=01f6

  IVT5508I DISPLAY ACCEPTED
  IVT5549I PROCESSING DISPLAY CSM COMMAND - OWNERID SPECIFIED
  IVT5530I BUFFER BUFFER
  IVT5551I SIZE   SOURCE             STORAGE ALLOCATED TO OWNER
  IVT5532I ----------------------------------------------------
  IVT5553I    4K  ECSA                              256K
```

```
IVT5554I TOTAL  ECSA                            256K
IVT5532I ------------------------------------------------------
IVT5553I    4K  DATA SPACE 64                   128K
IVT5554I TOTAL  DATA SPACE 64                   128K
IVT5532I ------------------------------------------------------
IVT5554I TOTAL  DATA SPACE                      128K
IVT5532I ------------------------------------------------------
IVT5556I TOTAL  FOR OWNERID                     384K
IVT5557I OWNERID: ASID = 01F6  JOBNAME = TCPIP
IVT5599I END
```

For additional information about the command, see z/OS Communications Server: SNA Operation.

Limiting TCP/IP common and private storage utilization

You can limit the amount of common and private storage that TCP/IP can use by coding the GLOBALCONFIG parameters ECSALIMIT and POOLLIMIT in the TCP/IP profile.

- The ECSALIMIT parameter specifies the maximum amount of common storage that TCP/IP can use.
- The POOLLIMIT parameter specifies the maximum amount of TCP/IP private storage that TCP/IP can use.

The ECSALIMIT parameter ensures that TCP/IP does not overuse common storage. It can improve system reliability by limiting TCP/IP's storage usage. The limit must account for peak storage usage during periods of high system activity or TCP/IP storage abends might occur. The limit does not include the CSM storage used by TCP/IP.

Tip: Care should be taken when coding the ECSALIMIT parameter. Setting it too low can cause TCP/IP to terminate prematurely.

Specifying a nonzero ECSALIMIT value enables warning messages EZZ4360I, EZZ4361I, and EZZ4362I to be issued when a storage shortage occurs.

If necessary, the ECSALIMIT and POOLLIMIT parameter values on the GLOBALCONFIG statement in the TCP/IP profile may be increased with a VARY TCPIP,,OBEYFILE command. For additional information regarding the VARY TCPIP,,OBEYFILE command, see z/OS Communications Server: IP System Administrator's Commands.

See z/OS Communications Server: IP Configuration Reference for more information regarding use of the GLOBALCONFIG statement ECSALIMIT and POOLLIMIT in the TCP/IP profile.

Limiting CSM storage utilization

CSM storage limits are located in the SYS1.PARMLIB member IVTPRM00. The values you can allocate are:

- ECSA MAX - the maximum amount of ECSA storage that CSM can allocate.
- FIXED MAX - the maximum amount of fixed storage that CSM can allocate. This includes fixed CSM ECSA, CSM data space storage, and high virtual common (HVCOMM) storage above the 2 GB bar.
- HVCOMM MAX - the maximum amount of HVCOMM storage that CSM can allocate.

If you do not specify values in the IVTPRM00 parmlib member, the system uses the default values of 100M ECSA, 200M FIXED, and 2000M HVCOMM. You can change these values dynamically with the MODIFY CSM command. If the limit specified by these values is reached, results are unpredictable. TCP/IP might not be able to continue. IVTxxxx messages will be issued if CSM is unable to obtain storage. See z/OS MVS Initialization and Tuning Reference for additional information about the IVTPRM00 parmlib member.

To change your CSM settings dynamically, issue the following command:

MODIFY net,CSM,ECSA=*value*,FIXED=*value*,HVCOMM=*value*

where the ECSA value is in the range 1024 KB - 2048 MB, the FIXED value is in the range 1024 KB - 30720 MB, and the HVCOMM value is in the range 100 MB - 999999 MB. Additional information regarding the MODIFY command for CSM can be found in z/OS Communications Server: SNA Operation.

Storage messages

The following messages are issued for TCP/IP common or private storage shortage problems. The messages are documented in greater detail in z/OS Communications Server: IP Messages Volume 4 (EZZ, SNM).

Common storage messages are as follows:

- EZZ4360I *jobname* ECSA CONSTRAINED
- EZZ4361I *jobname* ECSA CRITICAL
- EZZ4362I *jobname* ECSA EXHAUSTED
- EZZ4363I *jobname* ECSA SHORTAGE RELIEVED

Specifying a nonzero ECSALIMIT value enables warning messages EZZ4360I, EZZ4361I and EZZ4362I to be issued when a storage shortage occurs. Details regarding the TCP/IP profile GLOBALCONFIG parameter ECSALIMIT can be found in z/OS Communications Server: IP Configuration Reference.

Private storage messages are as follows:

- EZZ4364I *jobname* POOL CONSTAINED
- EZZ4365I *jobname* POOL CRITICAL
- EZZ4366I *jobname* POOL EXHAUSTED
- EZZ4367I *jobname* POOL SHORTAGE RELIEVED

Specifying a nonzero POOLLIMIT enables warning messages EZZ4364I, EZZ4365I, and EZZ4366I to be issued when a storage shortage occurs. Details regarding the TCP/IP profile, GLOBALCONFIG parameter POOLLIMIT, can be found in z/OS Communications Server: IP Configuration Reference. If storage limits were set using the TCP/IP profile GLOBALCONFIG ECSALIMIT statement or the GLOBALCONFIG POOLLIMIT statement look for the TCP/IP warning messages described previously that are issued each time a storage limit boundary is crossed. These messages might indicate a need to raise the limits.

CSM messages always start with the message prefix IVT. For a complete list of messages issued by CSM, see z/OS Communications Server: SNA Messages.

CSM messages identify whether the storage problem is related to CSM ECSA or CSM fixed storage. Examine the IVTPRM00 parmlib member to determine whether

the limits for the particular type of CSM storage that is depleted should be increased. Issue the Display CSM command to get more details on current CSM allocation and limits. As previously described, CSM limits can be increased using the Modify CSM command without reloading the initial program.

For more information about the Display CSM and Modify CSM commands, see z/OS Communications Server: SNA Operation.

Sysplex Problem Detection and Recovery (SPDR) storage messages

Critical storage shortages for CSM, ECSA, or PRIVATE are always detected by SPDR. Storage failures (when GLOBALCONFIG LIMITS values are not coded) are detected only when an allocation in the Sysplex code fails.

Sysplex Problem Detection and Recovery (SPDR) issues one of the following messages when a storage request for common, private, or CSM storage cannot be satisfied.

- EZD1170E
- EZD1187E
- EZZ9679E

EZD1170E
tcpstackname WAS NOT ABLE TO GET TCP/IP *storagetype* STORAGE

When the TCP/IP profile GLOBALCONFIG statement ECSALIMIT or POOLLIMIT parameter is not coded, message EZD1170E is issued. In this situation, the storage request fails because ECSA or private storage is exhausted.

EZD1187E
tcpstackname WAS NOT ABLE TO GET TCP/IP *storagetype* STORAGE

When TCP/IP profile GLOBALCONFIG statement ECSALIMIT or POOLLIMIT is coded, EZD1187E is issued. In this situation the storage request fails because ECSA or private storage is critical.

EZZ9679E
tcpstackname DETERMINED THAT CSM WAS CRITICAL FOR AT LEAST *timevalue* SECONDS

SPDR issues EZZ9679E when CSM storage problems are detected. This message is issued when CSM storage has been critical for the configured value specified on the GLOBALCONFIG SYSPLEXMONITOR TIMERSECS parameter or the default value of 60 seconds if the parameter is not specified.

See z/OS Communications Server: IP Messages Volume 2 (EZB, EZD) for debug information when these messages are issued. See z/OS Communications Server: IP Configuration Guide for more information about Sysplex Problem Detection and Recovery.

Abends

There are abends for each of the three types of storage problems:

Common storage
Common storage shortages typically result in the following abends:

- ABEND878 RC04 or RC08

- ABEND80A RC04 or RC08
- ABEND4C5 rsn xxxx2500

For common storage problems, determine which jobs or address spaces are using an excessive amount of storage. To determine the users of common storage, enable common storage tracking (CSA Tracker). For information about how to activate and review data provided by common storage tracking, see z/OS MVS Initialization and Tuning Guide. The storage totals for TCP/IP in the CSA tracker report does not reflect all the storage in use by TCP/IP. A number of TCP/IP getmain transactions are issued with the owner as SYSTEM. This storage is reported as OWNER = SYSTEM, and not OWNER = TCP/IP.

Contact the IBM Support Center for the owner of the storage that you think is causing the problem. See "Collecting documentation to submit to the IBM Support Center" on page 951 for more information.

Private storage

Private storage shortages typically result in the following abends:

- ABEND878 RC0C or RC10
- ABEND80A RC0C or RC10

If the problem is with TCP/IP private storage, submit a problem record with the IBM TCP/IP support team for dump analysis. See "Collecting documentation to submit to the IBM Support Center" on page 951 for more information.

CSM storage

For CSM storage problems, review the output from any monitoring you have been doing for CSM storage usage. Determine the largest users of the CSM ECSA and dataspace pools (4k, 16k, 32, 60k and 180k). See z/OS Communications Server: SNA Operation for additional information.

Contact the IBM Support Center for the owner of the CSM storage that you think is being used in excess. See "Collecting documentation to submit to the IBM Support Center" on page 951 for more information.

If you have not been tracking CSM storage utilization, see "Monitoring storage utilization" on page 942 to determine how to monitor this storage for use in problem diagnosis.

Diagnosing a storage problem

There are several causes of storage problems. When diagnosing a storage problem, whether it is an out of storage abend condition or a storage growth problem, you must first determine whether the storage problem is related to common, private, or CSM storage.

Steps for reviewing a storage problem

Determine whether a storage problem is related to common, private, or CSM storage.

Procedure

Perform the following steps:

1. Issue D TCPIP,,STOR and D Net,CSM,*Ownerid*=All commands to track storage usage.

2. Determine whether a lot of application data is accumulated on the queues for TCP connections. See “Steps for reviewing data on the queue if you are using IDS” on page 949.
3. Look in the system log for messages related to storage.
4. Run the EREP program against the SYS1.LOGREC log and review software records, looking for any storage-related abends.
5. Review the messages and abends you found to determine whether they indicate a common, private, or CSM storage problem.
6. Review your storage settings for any identified problem area (for example common, CSM, or private storage).
7. Compare your current storage usage to your previous usage. If usage has increased, do you know of any situations that would cause increased usage (for example, new applications that use common storage or increased connections)?
8. Review your response to step 6. Can the storage problem be resolved by increasing your storage limits?
9. If you have a dump that was created as a result of the storage problem, proceed to step 11.
10. If no dump was taken, see “Collecting documentation to submit to the IBM Support Center” on page 951 before proceeding to step 9. If you are unable to obtain a slip or console dump of the problem, a stand-alone dump may be the only method to gather documentation for the IBM Support Center. See z/OS MVS Diagnosis: Tools and Service Aids to determine how to take a stand-alone dump.
11. Call the IBM Support Center for assistance in reviewing the documentation.

Results

You know you are done when you have determined the storage problem can be remedied by increasing your storage limits or you have called the IBM Support Center for assistance.

Steps for reviewing data on the TCP queues

This topic provides the steps to be taken to determine whether data accumulated on the queues for TCP connections is contributing to a storage problem. The steps to be taken depend on whether you are using the intrusion detection services (IDS) and, if you are using IDS, how you have IDS configured.

Procedure

Complete the steps under the first bullet that describes your configuration:

- If you are using IDS, and you have IDS TCP Queue Size attack detection enabled, and your IDS policy for the TCP Queue Size attack type indicates that connections should be reset, then TCP connections with excessive or old data accumulated on their queues are automatically reset. Data accumulated on TCP queues should not be contributing to the storage problem. Proceed to Step 3 in “Steps for reviewing a storage problem” on page 946.
- Complete the following steps if you are using IDS, and you have IDS TCP Queue Size attack detection enabled, and your IDS policy for the TCP Queue Size attack type indicates that events should be logged to the system console:
  1. Look for console message groups that begin with message EZZ8761I and that indicate an event type of TCP Queue Constrained on message EZZ8762I.

This message group indicates that the TCP connection identified by messages EZZ8764I and EZZ8765I had excessive or old data accumulated on one of its queues. The following is an example of this message group:

```
EZZ8761I IDS EVENT DETECTED
EZZ8730I STACK TCPCS3
EZZ8762I EVENT TYPE: TCP QUEUE CONSTRAINED
EZZ8763I CORRELATOR 21 - PROBEID 040A0001
EZZ8764I SOURCE IP ADDRESS 4.4.4.4 - PORT 301
EZZ8765I DESTINATION IP ADDRESS 3.3.3.3 - PORT 300
EZZ8766I IDS RULE TCPQueSz
EZZ8767I IDS ACTION QueSzAction
```

The correlator value in message EZZ8763I identifies a unique instance of excessive or old data accumulated on a queue for a TCP connection. When the excessive or old data is removed from the queue, a similar message group that indicates an event type of TCP Queue Unconstrained on message EZZ8762I is issued. The same correlator value is used in the message group for this TCP Queue Unconstrained event as was used in the corresponding message group for the TCP Queue Constrained event.

- If none of these message groups have been issued, the data accumulated on TCP queues is not contributing to the storage problem. Proceed to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946.
- If these message groups have been issued, then go to Step 2.

2. For each TCP Queue Constrained message group located in step 1, look for a console message group that begins with message EZZ8761I and that indicates an event type of TCP Queue Unconstrained and includes the same correlator value as the TCP Queue Constrained message group. The following is an example of this message group:

```
EZZ8761I IDS EVENT DETECTED
EZZ8730I STACK TCPCS3
EZZ8762I EVENT TYPE: TCP QUEUE UNCONSTRAINED
EZZ8763I CORRELATOR 21 - PROBEID 040A0002
EZZ8764I SOURCE IP ADDRESS 4.4.4.4 - PORT 301
EZZ8765I DESTINATION IP ADDRESS 3.3.3.3 - PORT 300
EZZ8766I IDS RULE TCPQueSz
EZZ8767I IDS ACTION QueSzAction
```

   - If this message group has been issued, the excessive or old data that was on the queue for this TCP connection has been processed and is not contributing to the storage problem. Continue with the next TCP Queue Constrained message group located in step 1.
   - If this message group has not been issued, the excessive or old data is still on the queue for this TCP connection. Resolve the problem that is causing the data to not be processed, or reset the connection to release the storage.

3. If the storage problem is not resolved, proceed to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946.

- Complete the following steps if either you are using IDS, and you have IDS TCP Queue Size attack detection enabled, and your IDS policy for the TCP Queue Size attack type indicates that events should not be logged, or TRMD and syslogd are not both active:

  1. Issue the Netstat ALL/-A command to determine whether a lot of application data has accumulated on the queues for TCP connections.
     - If no connections are found with a significant amount of data on their queues, then data accumulated on TCP queues is not contributing to the storage problem. Proceed to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946.

- If connections are found with a significant amount of data on their queues, then go to Step 2.

2. For each connection with a significant amount of data on any of its queues, resolve the problem that is causing the data to not be processed, or reset the connection to release the storage.
3. If the storage problem is not resolved, proceed to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946.

- Complete the following steps if none of the previous bullets describes your configuration:
  1. Look in the syslogd output for messages EZZ8662I, EZZ8664I, or EZZ8666I. These messages indicate that excessive or old data is accumulating on the receive, send, or out-of-order queue for a TCP connection. The correlator value in these messages identifies a unique instance of excessive or old data accumulated on a queue for a TCP connection. This same correlator value will appear in a EZZ8663I, EZZ8665I, or EZZ8667I message that is issued when the excessive or old data has been removed from the TCP queue.
     - If EZZ8662I, EZZ8664I, and EZZ8666I messages do not appear in syslogd, then data accumulated on TCP queues is not contributing to the storage problem. Proceed to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946.
     - If EZZ8662I, EZZ8664I, and EZZ8666I messages do appear in syslogd, then go to Step 2.
  2. For each EZZ8662I, EZZ8664I, or EZZ8666I message located in step 1, look for a corresponding EZZ8663I, EZZ8665I, or EZZ8667I message that includes the same correlator.
     - If a corresponding message with the same correlator has been issued, the excessive or old data that was on the queue for this TCP connection has been processed and is not contributing to the storage problem. Continue with the next EZZ8662I, EZZ8664I, or EZZ8666I message located in step 1.
     - If a corresponding message with the same correlator has not been issued, the excessive or old data is still on the queue for this TCP connection. Resolve the problem that is causing the data to not be processed, or reset the connection to release the storage.
  3. If the storage problem is not resolved, proceed to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946.

Steps for reviewing data on the queue if you are using IDS

If you are using IDS, review the data on the queue.

Before you begin

You need to have issued the D TCPIP,,STOR and D Net,CSM,*Ownerid*=All commands to track storage use.

Procedure

1. If TCP QUEUE Size attack detection is enabled with console logging, look for console message groups that begin with message EZZ8761I and that indicate a TCP QUEUE Constrained event.
   -

If these message groups have not been issued, then data accumulated on TCP queues is not contributing to the storage problem and you can proceed to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946.

- If these message groups have been issued, then go to Step 2.

2. Depending on your configuration, take the following actions to determine whether data is accumulated on the queue:

If . . .	Then . . .	And take the following actions...
TRMD or syslog are not active, or TCP QUEUE Size attack detection is not configured to log to syslogd	Issue the Netstat ALL/-A command to determine whether a lot of application data has accumulated on the queues for TCP connections.	• If no data, or an insignificant amount of data, has accumulated on the queue, go to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946. • If data has accumulated on the queue, then resolve the problem that is causing the data to not be processed, or reset those connections to release the storage.
TRMD and syslog are active, and TCP QUEUE Size attack detection is configured to log to syslogd	Look in the syslogd output for messages EZZ8621I, EZZ86641I, or EZZ8666I. These messages indicate that excessive or old data is accumulating on the receive, send, or out-of-order queue for a TCP connection.	• If these messages do not appear in syslogd, go to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946 • If one or more of these messages do appear in syslogd, determine whether there is a corresponding message (EZZ863I, EZZ8665I, or EZZ8667I) that indicates that the accumulated data has been processed. If no corresponding message appears, resolve the problem that is causing the data to not be processed, or reset those connections to release the storage.

Results

You can now perform the steps for the decision you have made.

Steps for reviewing data on the queue if you are not using IDS

Before you begin, you need to have issued the D TCPIP,,STOR and D Net,CSM,*Ownerid*=All commands to track storage use.

Depending on your configuration, take the following actions to determine whether data is accumulated on the queue:

If . . .	Then . . .	And take the following actions...
TCP QUEUE Size attack detection is not configured to log to syslogd	Issue the Netstat ALL/-A command to determine whether a lot of application data has accumulated on the queues for TCP connections.	• If no data, or an insignificant amount of data, has accumulated on the queue, go to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946. • If data has accumulated on the queue, then resolve the problem that is causing the data to not be processed, or reset those connections to release the storage.
TCP QUEUE Size attack detection is configured to log to syslogd	Look in the syslogd output for messages EZZ8621I, EZZ86641I, or EZZ8666I. These messages indicate that excessive or old data is accumulating on the receive, send, or out-of-order queue for a TCP connection.	• If these messages do not appear in syslogd, go to Step 3 on page 947 in "Steps for reviewing a storage problem" on page 946 • If one or more of these messages do appear in syslogd, determine whether there is a corresponding message (EZZ863I, EZZ8665I, or EZZ8667I) that indicates that the accumulated data has been processed. If no corresponding message appears, resolve the problem that is causing the data to not be processed, or reset those connections to release the storage.

You can now perform the steps for the decision you have made.

Collecting documentation to submit to the IBM Support Center

If your storage problem is caused by TCP/IP common or private storage usage, ensure that you have a dump of TCP/IP; IBM service will need to review it. If you did not get a system dump for the abend, or if you want to obtain a dump of TCP/IP to perform a storage analysis, use the following table for commands you can issue.

Table 106. Commands for various types of dumps

Commands	Type of dump

`SL SET,COMP=xxx,ACTION=SVCD,JOBNAME=tcpipprocname,`
`SDATA=(NUC,RGN,CSA,SQA,LSQA,TRT),END`
 	
where *xxx* equals the abend code you are receiving (for example, 878).
	SLIP DUMP of the abend

Table 106. Commands for various types of dumps (continued)

Commands	Type of dump
SL SET,COMP=*xxx*,ACTION=SVCD,JOBNAME=*tcpipprocname*, DSPNAME=(1.CSM*), SDATA=(NUC,RGN,CSA,SQA,LSQA,TRT),END where *xxx* equals the abend code you are receiving (for example, 878).	SLIP dump of the abend if your storage problem is related to TCP/IP's CSM storage usage (includes the CSM data spaces in your dump).
DUMP COMM=(*tcpip storage growth*) R *xx*,JOBNAME=(*tcpipprocname*),CONT R *xx*,SDATA=(NUC,RGN,CSA,SQA,LSQA,TRT),END	Console dump of TCP/IP
SL SET,MSGID=*zzzzz*,ACTION=SVCD,JOBLIST=(*tcpipprocname*, VTAM address space name),DSPNAME=(1.CSM*), SDATA=(NUC,RGN,CSA,SQA,LSQA,TRT),END where *zzzzz* causes a dump when the message identified in the slip is issued (that is, CSM message IVT5562I).	MSGID slip can be used to take a dump when a particular message is issued.
DUMP COMM=(*tcpip storage growth*) R *xx*,JOBNAME=(*tcpipprocname*),CONT R *xx*,SDATA=(NUC,RGN,CSA,SQA,LSQA,TRT),CONT R *xx*,DSPNAME=(1.CSM*),END	Console dump when the storage problem is related to CSM storage usage (includes CSM data spaces).
DUMP COMM=('*storage growth*') R *xx*,JOBNAME=(*tcpipprocname*,*applname*),CONT R *xx*,SDATA=(NUC,RGN,CSA,SQA,LSQA,LPA,TRT),END	Dump of TCP/IP and any TCP/IP related applications that may be having storage problems (for example Omproute, FTP)

Note: Wildcards (*) allow you to use a single specification to indicate a number of address spaces whose names match the wildcard pattern. This can be useful if you need to dump multiple TCP/IP stacks. You can specify a wildcard on the JOBLIST and DSPNAME parameters of a SLIP. And, the JOBNAME and DSPNAME parameters of the console dump command. For information about how to use wildcards in a SLIP and DUMP command, see z/OS MVS System Commands.

Submit the console log, and all dumps to the IBM Support Center for review.

Part 4. Appendixes

© Copyright IBM Corp. 2000, 2015

Appendix A. Overview of internetworking

This topic gives an overview of internetworking and contains the following subtopics:

- "Maximum transmission unit (MTU)" on page 956
- "Fiber Distributed Data Interface (FDDI)" on page 957
- "Token-Ring IEEE 802.5" on page 958
- "IEEE 802.3" on page 959
- "Ethernet — DIX V2" on page 959
- "Subnetwork Access Protocol (SNAP)" on page 960
- "IP routing" on page 961
- "Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6)" on page 961
- "Direct routing" on page 964
- "Indirect routing" on page 965
- "Simplified IP datagram routing algorithm" on page 965
- "IPv4 subnetting" on page 966
- "IPv6 prefixes" on page 967
- "Simplified IP datagram routing algorithm with subnets" on page 967
- "Static routing" on page 969
- "Dynamic routing" on page 969

Networking with TCP/IP connects different networks so that they form one logical interconnected network. This large overall network is called an *internetwork*, or more commonly, an *intranet* or *internet*. Each network uses its own physical layer, and the different networks are connected to each other by means of machines that are called *gateways*.

Gateways transfer IP datagrams between networks. This function is called *routing*; therefore, the internet gateways are often called *routers*. Within this appendix, the terms router and gateway are synonymous; both refer to a machine that transfers IP datagrams between different networks.

If IP datagrams are not passed properly over a bridge, none of the higher TCP/IP protocols or applications work correctly. For a discussion of bridges, see *TCP/IP Tutorial and Technical Overview*.

Linking networks in this way takes place at the network level of the International Organization for Standardization (ISO). It is possible to link networks at a lower-level layer using *bridges*. Bridges link networks at the ISO data link layer. Bridges pass packets or frames between different physical networks regardless of the protocols contained within them. An example of a bridge is the IBM 8209, which can interconnect an Ethernet network and a token-ring network.

A bridge does not connect TCP/IP networks together. It connects physical networks together that still forms the same TCP/IP network. (A bridge does not do IP routing.)

© Copyright IBM Corp. 2000, 2015

Figure 109 depicts a router and a bridge. The router connects Network 1 to Network 2 to form an intranet.

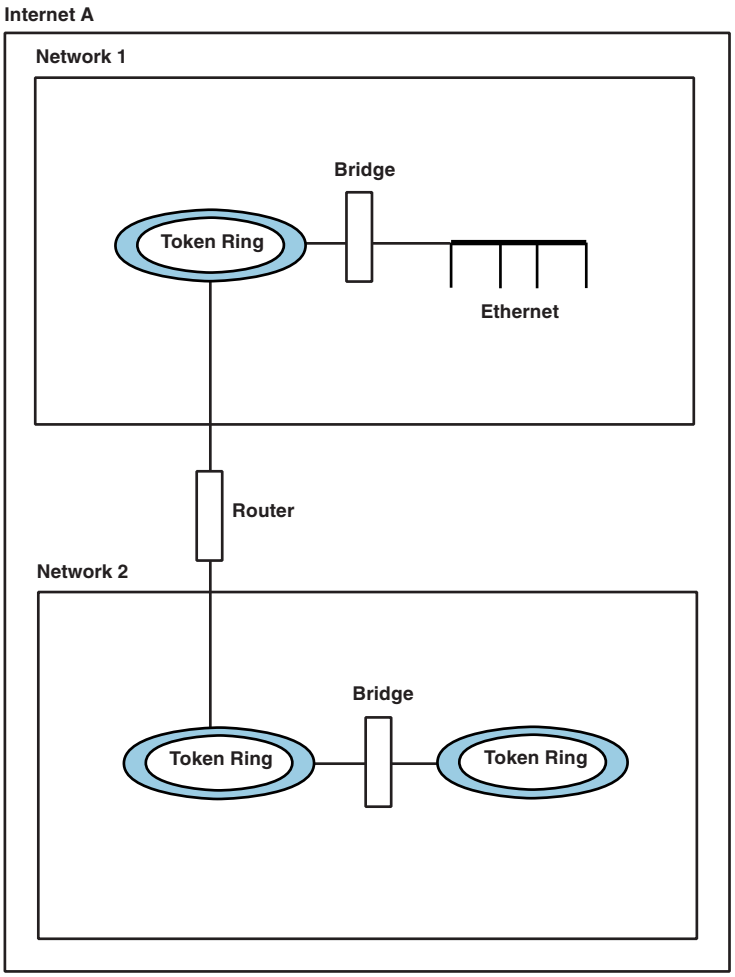

Figure 109. Routers and bridges within an internet

Maximum transmission unit (MTU)

Different physical networks have different maximum frame sizes. Within the different frames, there is a maximum size for the data field. This value is called the *maximum transmission unit* (MTU), or maximum packet size in TCP/IP terms.

Figure 110 on page 957 shows the relationship between MTU and frame size.

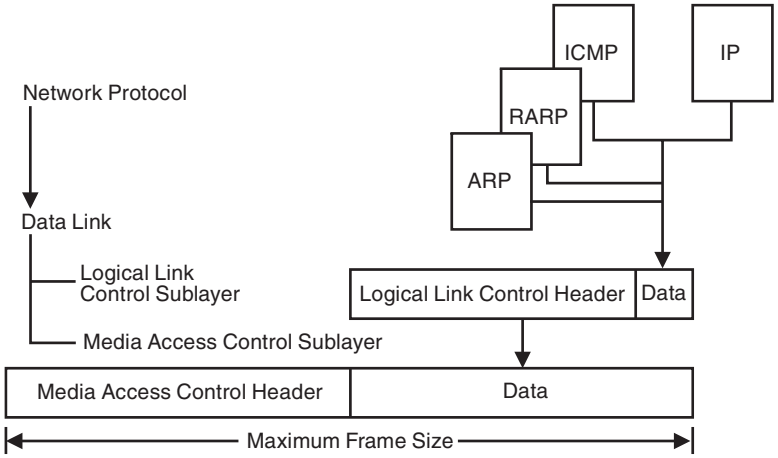

Figure 110. Relationship of MTU to frame size

If an IP datagram is to be sent out onto the network and the size of the datagram is bigger than the MTU, IP fragments the datagram into multiple fragments, so that it fits within the data fields of the frames. If the MTU is larger than the network can support, then the data is lost.

The value of MTU is especially important when bridging is used because of the different network limits. RFC 791 —Internet Protocols states that all IP hosts must be prepared to accept datagrams of up to 576 bytes.

The minimum MTU for IPv6 is 1280. See RFC 2460, Internet Protocol, Version 6 (IPv6) Specification for more information.

You can configure an MTU value by using the MTU parameter on the BEGINROUTES statement.

Fiber Distributed Data Interface (FDDI)

The FDDI specifications define a family of standards for 100 Mbps fiber optic LANs that provide the physical layers and media access control sublayer of the data link layer, as defined by the ISO/OSI Model.

IP-FDDI defines the encapsulating of IP datagrams and ARP requests and replies in FDDI frames.

All frames are transmitted in standard IEEE 802.2 LLC Type 1 Unnumbered Information format, with the DSAP and SSAP fields of the 802.2 header set to the assigned global SAP value for SNAP (decimal 170). The 24-bit Organization Code in the SNAP header is set to zero, and the remaining 16 bits are the EtherType from Assigned Numbers:

- 2048 for IP
- 2054 for ARP

Typically, the MTU is set to 4352.

Mapping of 32-bit internet addresses to 48-bit FDDI addresses is done by the ARP dynamic discovery procedure. The broadcast internet addresses (whose <host address> is set to all ones) are mapped to the broadcast FDDI addresses (all ones).

IP datagrams are transmitted as a series of 8-bit bytes using the usual TCP/IP transmission order called *big-endian* or *network byte order.*

For more information about FDDI architecture, see *LAN Concepts and Products*.

Token-Ring IEEE 802.5

When a token-ring frame passes through a bridge, the bridge adds information to the routing information field (RIF) of the frame (assuming that the bridge supports source route bridging). The RIF contains information concerning the route taken by the frame and, more importantly, the maximum amount of data that the frame can contain within its data field. This is called the maximum information field (I-field). The value specified for the maximum I-field is sometimes referred to as the largest frame size, but this means the largest frame size, excluding headers. See Figure 111 for details on the relationship of the I-field to the header fields.

Guideline: It is important to be aware that the IBM implementation limits the number of bridges through which a frame can be passed to seven. An attempt to pass a frame through an eighth bridge fails.

The maximum I-field is always decreased by a bridge when it cannot handle the value specified. So, for a given path through a number of token-ring bridges, the maximum I-field is the largest value that all of the bridges support. This value is specified in the Routing Control (RC) field within the RIF as shown in Figure 111.

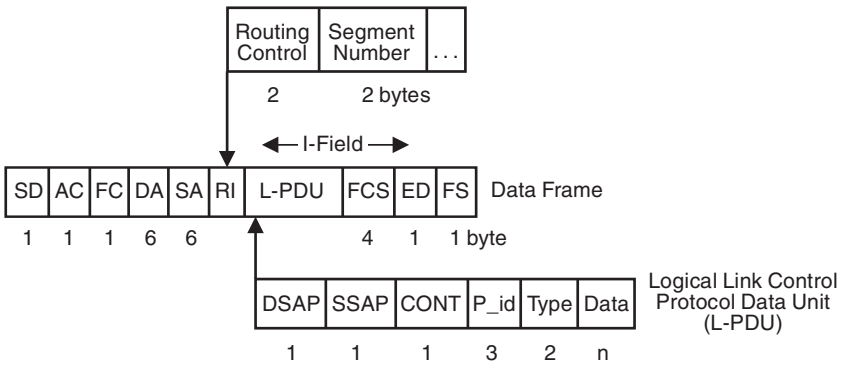

Figure 111. Format of an IEEE 802.5 token-ring frame

The size of the MTU is the maximum amount of data that is allowed within a frame. The token-ring architecture specifies the maximum value of the I-field in the data frame, which corresponds to the maximum size of the L-PDU. The maximum I-field value is determined by the bit configuration in the RC field, and is present in all routed frames.

Table 107 on page 959 shows the relationship between the RC field and the maximum I-field values.

Table 107. Relationship between RC field and maximum I-field value

Routing control field	Maximum I-field in bytes
x000 xxxx xxxx xxxx	516
x001 xxxx xxxx xxxx	1500
x010 xxxx xxxx xxxx	2052
x011 xxxx xxxx xxxx	4472
x100 xxxx xxxx xxxx	8144
x101 xxxx xxxx xxxx	11407
x110 xxxx xxxx xxxx	17800

Figure 111 on page 958 shows that, within the L-PDU, the Logical Link Control (LLC) header uses 8 bytes. Thus the MTU value is 8 bytes less than the maximum I-field. The L-PDU contains a SNAP header, as described in "Subnetwork Access Protocol (SNAP)" on page 960. Follow this example to calculate the MTU for a token-ring. The token-ring bridges always adjust the value of the maximum I-field to that of the smallest one in the path. Ensure that the MTU value is less than the value specified by the bridge.

Typically, within a 4-Mbps token-ring network, the value of maximum I-field is 2052 bytes. Therefore, the MTU would be set to 2044 bytes (2052 minus 8 bytes for the LLC header).

IEEE 802.3

The frame used in IEEE 802.3 Ethernet networks is shown in Figure 112.

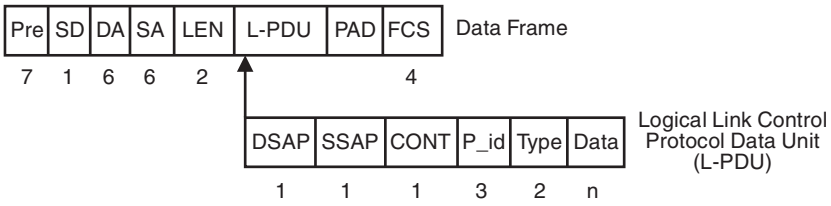

Figure 112. Format of an IEEE 802.3 frame

The maximum size of the L-PDU for a 10Mbps network is 1500 bytes. Because 8 bytes are used within the L-PDU for the LLC header, this means that the maximum size of the data field is 1492 bytes. Therefore, set the MTU for IEEE 802.3 networks to 1492 bytes.

Ethernet — DIX V2

The frame used in DIX Ethernet networks is shown in Figure 113 on page 960.

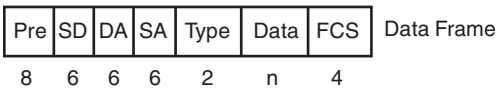

Figure 113. Format of an Ethernet V2 frame

There is no LLC data in an Ethernet V2 frame. The maximum size for the frame is 1526 bytes. This means that the data field can be 1500 bytes maximum. The MTU for Ethernet V2 can be set to 1500 bytes.

It is possible to bridge Ethernet V2 frames to either IEEE 802.3 or IEEE 802.5 networks; an LLC header is added or removed from the frame, as required, as part of the conversion when bridging.

Subnetwork Access Protocol (SNAP)

The TCP/IP software provides protocol support down to the ISO network layer. Following this layer is the data link layer, which can be separated into two sublayers. These are the *Logical Link Control* (LLC) and the *Media Access Control* (MAC) layers.

The IEEE 802.2 standard defines the LLC sublayer, and the MAC sublayer is defined in IEEE 802.3, IEEE 802.4, and IEEE 802.5.

The format of an IEEE 802.2 LLC header with the SNAP header is shown in Figure 114.

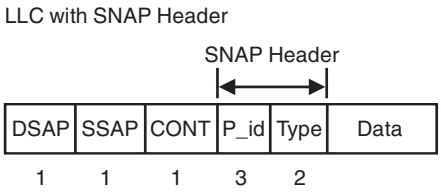

Figure 114. SNAP header

The values of the fields in the LLC header when a SNAP header is used are specified in RFC 1042 - Standard for Transmission of IP Datagrams over IEEE 802 Networks. The values specified are:

Field **Value**

DSAP X'AA'

SSAP X'AA'

CONT
X'03' Specifies unnumbered information (UI)

P_id X'00 00 00'

Type

X'8006' — ARP
X'8035' — RARP
X'86dd' — IPv6

IP routing

IP routing is based on routing tables held within a router or internet host. These tables contain routes which can either be *static* or *dynamic*. Typically, static routes are predefined within a configuration file, and dynamic routes are *learned* from the network, using a *routing* protocol.

Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6)

There are two Internet protocols used to assign addresses to links on a host, Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6). The majority of current internets use IPv4. This protocol is nearly 20 years old and is approaching the limits of the node addresses that its 32 bit addresses allow. IPv6 is the next generation of the Internet Protocol, designed to replace IPv4. Among other advantages, the 128 bit addresses defined by IPv6 provide nearly limitless addresses.

Although IPv6 is expected to eventually replace IPv4, they are likely to coexist for a number of years during the transition.

Figure 115 on page 962 shows a simple network with a bridge and a router.

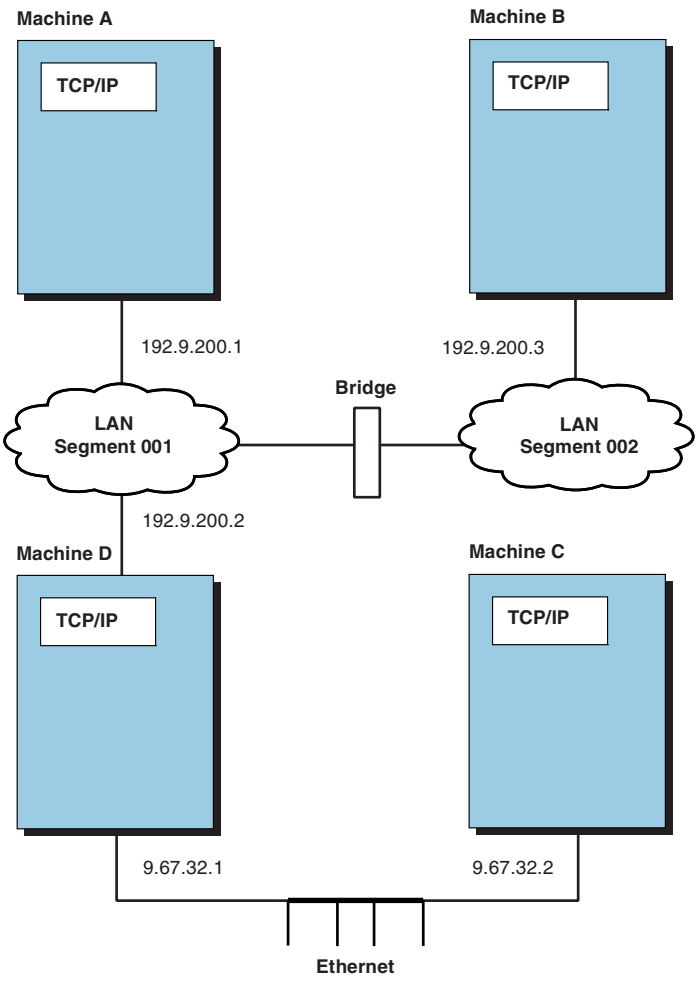

Figure 115. Routing and bridging

Machine D is acting as an IP router and transfers IP datagrams between the class C, 192.9.200, network and the class A, 9.67.32 network. It is important to note that for Machine B to communicate with Machine C using TCP/IP, both Machine D and the bridge have to be correctly configured and working.

TCP/IP uses the HOME statements, defined in the data set *hlq*.PROFILE.TCPIP, to assign home addresses and associated link names. HOME statements can be updated using the VARY TCPIP command. See the z/OS Communications Server: IP Configuration Reference for more information about both the HOME statements.

Internet Protocol Version 4 (IPv4)

A link on a host on an intranet is identified by its *IP address*. *Internet Protocol* (IP) is the protocol that is used to deliver datagrams between such hosts. It is assumed the reader is familiar with the TCP/IP protocols. Details of some of the protocols can be found in the *TCP/IP Tutorial and Technical Overview*. Specific information relating to the Internet Protocol can be found in RFC 791.

An IPv4 address is a 32-bit address that is usually represented in dotted decimal notation, with a decimal value representing each of the four octets (bytes) that make up the address. For example:

```
00001001010000110110000100000010        32-bit address
00001001 01000011 01100001 00000010     4 octets
   9        67       97       2         dotted decimal notation (9.67.97.2)
```

The IPv4 address consists of a *network address* and a *host address*. Within the Internet, the network addresses are assigned by a central authority, the *Network Information Center* (NIC). The portion of the IPv4 address that is used for each of these addresses is determined by the class of address. There are three commonly used classes of IPv4 addresses (see Figure 116).

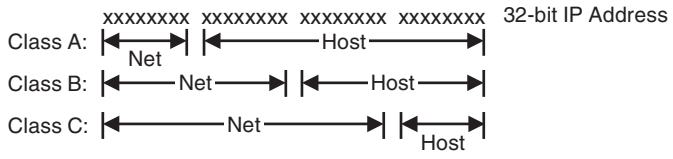

Figure 116. Classes of IPv4 addresses

The class of the address is determined by the first octet of the IPv4 address. Figure 117 shows how the class of address is determined. The figure also shows Class D addresses. Class D addresses represent multicast groups, not network IP addresses. Multicast group addresses consist of the high-order, four bits of 1110 and the remaining 28 bits, which form a multicast group ID.

```
32-bit address          xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx

Class A                 0xxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx
               min      00000000
               max      01111111
               range    1 - 126    (decimal notation; 0 and 127 are reserved)

Class B                 10xxxxxx xxxxxxxx xxxxxxxx xxxxxxxx
               min      10000000
               max      10111111
               range    128 - 191  (decimal notation)

Class C                 110xxxxx xxxxxxxx xxxxxxxx xxxxxxxx
               min      11000000
               max      11011111
               range    192 - 223  (decimal notation)

Class D                 1110xxxx xxxxxxxx xxxxxxxx xxxxxxxx
               min      11100000
               max      11101111
               range    224-239 (decimal notation)
```

Figure 117. Determining the class of an IPv4 address

As shown in Figure 117, the value of the bits in the first octet determine the class of address, and the class of address determines the range of values for the network and host segment of the IPv4 address. For example, the IPv4 address 9.67.97.2 would be a class A address, since the first two bits in the first octet contain B'00'. The network part of the IPv4 address is 9 and the host part of the IPv4 address is 67.97.2.

See RFC 1166–Internet Numbers for more information about IPv4 addresses. See RFC 1060–Assigned Numbers for more information about reserved network and host IPv4 addresses, such as a *network broadcast address*.

Internet Protocol Version 6 (IPv6)

As described above, 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.

There are three conventional forms for representing IPv6 addresses as text strings:

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

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

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. The following is the only exception.

It is common in some styles of IPv6 addresses to contain long strings of zero bits. To make writing addresses containing zero bits easier, a special syntax is available to compress the zeros. Use two colons (::) to indicate multiple groups of 16 bits of zeros. The two colons (::) can appear only once in an address. The two colons (::) can also be used to compress the leading zeros, the trailing zeros, or both in an address.

For example, the following addresses:

```
1080:0:0:0:8:800:200C:417A  a unicast address
FF01:0:0:0:0:0:0:101        a multicast address
0:0:0:0:0:0:0:1             the loopback address
0:0:0:0:0:0:0:0             the unspecified addresses
```

can be represented as:

```
1080::8:800:200C:417A       a unicast address
FF01::101                   a multicast address
::1                         the loopback address
::                          the unspecified addresses
```

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:x:d.d.d.d, where x is the hexadecimal value of the six high-order 16-bit pieces of the address, and d is the decimal value of the four low-order 8-bit pieces of the address (standard IPv4 representation). For example, `0:0:0:0:0:0:13.1.68.3` can be expressed in condensed form as `::13.1.68.3`

Direct routing

Direct routing can take place when two hosts are directly connected to the same physical network. This can be a bridged token-ring network, a bridged Ethernet, or a bridged token-ring network and Ethernet. The distinction between direct routing and indirect routing is that, with direct routing, an IP datagram can be delivered to the remote host without subsequent interpretation of the IP address, by an intermediate host or router.

In Figure 115 on page 962, a datagram traveling from Machine A to Machine B would be using direct routing, although it would be traveling through a bridge.

Indirect routing

Indirect routing takes place when the destination is not on a directly attached IP network, forcing the sender to forward the datagram to a router for delivery.

In Figure 115 on page 962, a datagram from Machine A being delivered to Machine C would be using indirect routing, with Machine D acting as the router (or gateway).

Simplified IP datagram routing algorithm

To route an IP datagram on the network, the algorithm shown in Figure 118 is used.

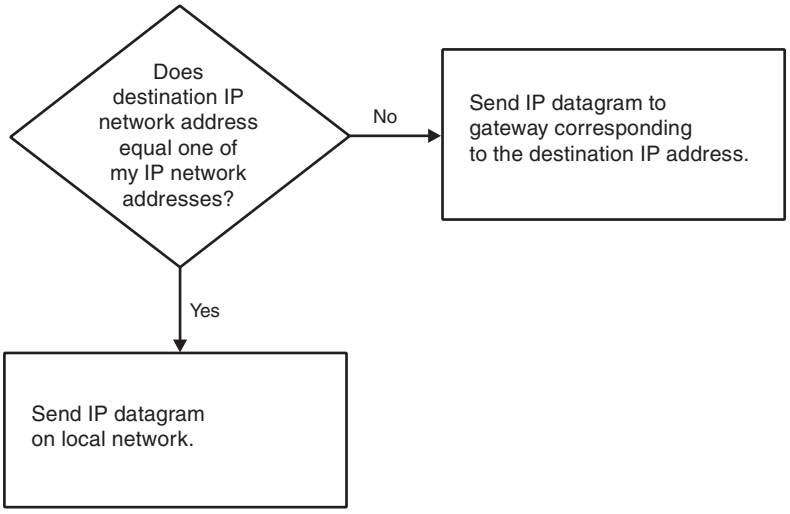

Figure 118. General IP routing algorithm

Using this general routing algorithm, it is easy to determine where an IP datagram is routed. Following is a simple example based on the configuration shown in Figure 115 on page 962.

```
Machine A   IP Address = 192.9.200.1

Routing Table

Destination     Gateway

192.9.200.1     192.9.200.1   (Machine A's network interface)

9.0.0.0         192.9.200.2   (Route to the 9.n.n.n address is
                               via Machine D, 192.9.200.2)
```

Machine A sends a datagram to host 192.9.200.3 (Machine B), using the direct route, 192.9.200.1 (its own network interface). Machine A sends a datagram to host 9.67.32.2 (Machine C), using the indirect route, 192.9.200.2 (Machine D), and Machine D then forwards the datagram to Machine C.

IPv4 subnetting

IPv4 allows for a variation of the network and host segments of an IP address, known as *subnetting*, can be used to physically and logically design a network. For example, an organization can have a single internet network address (NETID) that is known to users outside the organization, yet configure its internal network into different departmental subnets. Subnetwork addresses enhance local routing capabilities, while reducing the number of network addresses required.

To illustrate this, consider a simple example. Assume that we have an assigned class C network address of 192.9.200 for our site. This would mean that we could have host addresses from 192.9.200.1 to 192.9.200.254. If we did not use subnetting, then we could only implement a single IP network with 254 hosts. To split our site into two logical subnetworks, we could implement the network scheme shown in Figure 119:

```
Without Subnetting:
                                        Network        Host Address
                                        Address        Range
       192      9        200       host
    11000000 00001001 11001000 xxxxxxxx   192.9.200      1 - 254


With Subnetting:
                                        Subnet         Host Address  Subnet
                                        Address        Range         Value
       192      9        200   64  host
    11000000 00001001 11001000 01xxxxxx   192.9.200.64   65 - 126      01


                                        Subnet         Host Address  Subnet
                                        Address        Range         Value
       192      9        200   128 host
    11000000 00001001 11001000 10xxxxxx   192.9.200.128  129 - 190     10

The subnet mask would be

       255      255      255     192
    11111111 11111111 11111111 11000000
```

Figure 119. Subnetting scheme

z/OS TCP/IP uses a slightly different scheme for the subnet mask when defining the BEGINROUTES statements in the *hlq*.PROFILE.TCPIP data set and for displaying the subnet mask within a Netstat ROUTE/-r command. The subnet mask is applied only to the host segment of the IP address, and Netstat ROUTE/-r displays the subnet mask for only the host segment of the IP address. The subnet mask in the preceding chart as defined for z/OS TCP/IP would be:

```
   0        0        0       192       0.0.0.192
00000000 00000000 00000000 11000000
```

Although z/OS TCP/IP defines the subnet mask differently, the application of the subnet mask and subnet value to the IP address is consistent with RFC-architected routing algorithms. A subnet mask of 255 is used for the remainder of this section, to retain symmetry with other routing documents that use 255 as the subnet value for the network segment of an IP address.

Because subnets B'00' and B'11' are both reserved, only two subnets are available. All 0s and all 1s have a special significance in internet addressing and should be

used with care. Also notice that the total number of host addresses that we can use is reduced for the same reason. For instance, we cannot have a host address of 16 because this would mean that the subnet/host segment of the address would be B'0001000', which with the subnet mask we are using, would mean a subnet value of B'00', which is reserved.

The same is true for the host segment of the fourth octet. A fourth octet value of B'01111111' is reserved because, although the subnet of B'01' is valid, the host value of B'1' is reserved.

The network segment of the subnet mask is always assumed to be one, so each octet has a decimal value of 255. For example, with a class B address, the first two octets are assumed to be 255.255.

IPv6 prefixes

The IPv6 prefix concept is similar to IPv4 subnetting. An IPv6 address with a prefix is written as an IPv6 address followed by a decimal number representing the number of bits in the address that constitute the prefix. It is written as:

```
ipv6-address/prefix-length
```

where:

ipv6-address
is an IPv6 address in any notation

prefix-length
is 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 12AB00000000CD3 (hexadecimal):

```
12AB:0000:0000:CD30:0000:0000:0000:0000/60
12AB::CD30:0:0:0:0/60
12AB:0:0:CD30::/60
```

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

The node address

```
12AB:0:0:CD30:123:4567:89AB:CDEF
```

and its subnet number

```
12AB:0:0:CD30::/60
```

can be abbreviated as

```
12AB:0:0:CD30:123:4567:89AB:CDEF/60
```

Simplified IP datagram routing algorithm with subnets

When subnetting is used, the algorithm required to find a route for an IP datagram is similar to the one for general routing, with the exception that the addresses that are compared are the result of a logical AND of the subnet mask and the IP address.

For example:

```
IP address:       9.67.32.18   00001001 01000011 00100000 00010010
                                                      <AND>
Subnet Mask: 255.255.255.240   11111111 11111111 11111111 11110000

    Result of
    Logical AND:      9.67.32.16   00001001 01000011 00100000 00010000
```

The subnet address is 9.67.32.16, and it is this value that is used to determine the route used.

Figure 120 shows the routing algorithm used with subnets.

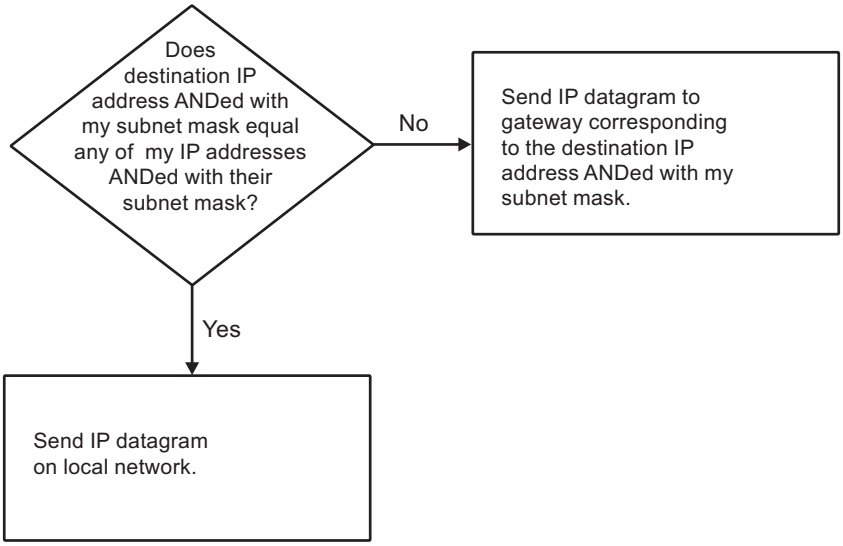

Figure 120. Routing algorithm with subnets

Figure 121 on page 969 shows how a subnet route is resolved.

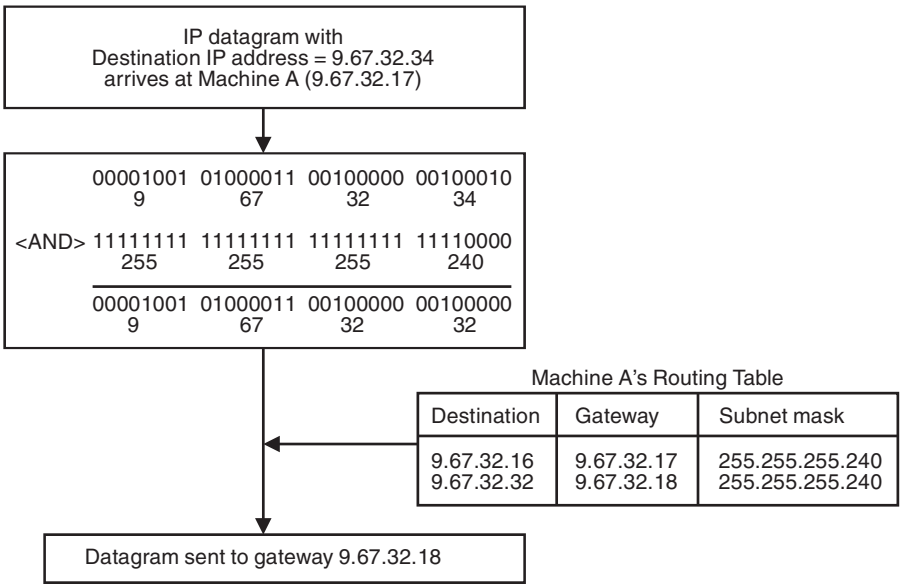

Figure 121. Example of resolving a subnet route

Static routing

Static routing, as the name implies, is defined within the local host, and must be manually changed as the network changes. Typically, a configuration file contains the definitions for directly-attached networks, routes for specific hosts, and a possible default route that directs packets to a destination for networks that are not previously defined.

Static routes can be defined using the z/OS TCP/IP BEGINROUTES statement to configure the internal routing tables; the statement is defined in the *hlq*.PROFILE.TCPIP data set. The internal routing tables for z/OS TCP/IP can be modified by one of the following ways:

- Changing the BEGINROUTES statement and recycling the TCP/IP address space.
- Using the VARY TCPIP,,OBEYFILE command.

See z/OS Communications Server: IP System Administrator's Commands for details about defining the BEGINROUTES statement.

Tip: When the BEGINROUTES statement is updated by using VARY TCPIP,,OBEYFILE, all previously defined static routes are discarded and replaced by the new BEGINROUTES definitions.

Dynamic routing

Dynamic routing is the opposite of static routing. A TCP/IP protocol is used to dynamically update the internal routing tables when changes to the network occur.

IPv4

For IPv4, there are two dynamic routing protocols available. One routing protocol is the Routing Information Protocol (RIP). It is implemented by the OMPROUTE routing applications. A newer protocol is open shortest path first (OSPF). It is implemented by OMPROUTE only. For more details about OMPROUTE, see Chapter 31, "Diagnosing OMPROUTE problems," on page 759. For configuration information about both applications, see z/OS Communications Server: IP Configuration Reference.

IPv6

For IPv6, dynamic routing is performed by the Router Discovery protocol and by the IPv6 OSPF and IPv6 RIP dynamic routing protocols of OMPROUTE. For more information about IPv6 dynamic routing, see z/OS Communications Server: IP Configuration Guide.

Appendix B. IKE protocol details

This topic provides an overview of the IKE protocols, IKE version 1 (IKEv1) and IKE version 2 (IKEv2).

IKE version 1 protocol

Overview of negotiating IKEv1 security associations

The ISAKMP protocol is a framework for dynamically establishing security associations and cryptographic keys in an Internet environment. This framework defines a set of message flows (exchanges) and message formats (payloads). ISAKMP defines a generic payload for key exchange information. This enables the ISAKMP protocol to manage cryptographic keys independent of the key exchange protocol that is used to generate them.

ISAKMP defers the interpretation of the key exchange payload to individual key exchange protocols. Internet Key Exchange (IKE) is such a protocol. IKE augments the ISAKMP protocol to facilitate the creation of authenticated keying material. IKE defines how keying material is generated. The exchanges that are defined by ISAKMP require authentication to take place, but they do not specify how authentication is to be performed. IKE defines how authentication is to be performed.

ISAKMP defines two phases of negotiation. Both of these phases are also applicable to the IKE protocol. The first phase is referred to as phase 1. In phase 1, two ISAKMP servers agree on how to protect traffic between themselves. This agreement results in the creation of an ISAKMP security association. The second phase is referred to as phase 2. In phase 2, security associations for other security protocols are established; for example, AH or ESP. Negotiations during each phase are accomplished using an ISAKMP-defined exchange or by an exchange that is specific to a key exchange protocol.

Phase 1

IKE supports two types of phase 1 exchanges:

- Main mode
- Aggressive mode

Both of these exchange modes are based on exchanges that are defined by ISAKMP. Main mode is an implementation of ISAKMP's Identity Protect exchange. Aggressive mode is an implementation of ISAKMP's Aggressive exchange.

IKE defines four techniques for authentication of phase 1 exchanges:

- Pre-shared key
- Signature-based
- Public key encryption
- Revised public key encryption

Restriction: Of these techniques, the z/OS IKE daemon supports only pre-shared key authentication and signature-based authentication using RSA signatures.

© Copyright IBM Corp. 2000, 2015

Main mode

A Main mode exchange is composed of six messages as shown in Figure 122.

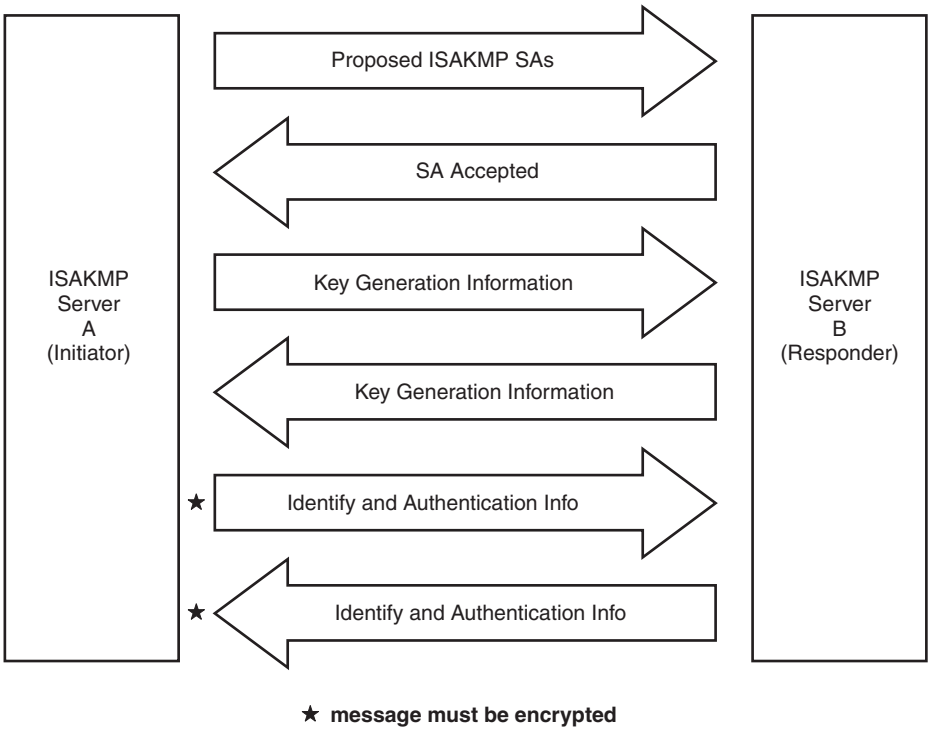

Figure 122. Main mode exchange

Messages 1 and 2 provide agreement on the negotiable attributes of the ISAKMP security association. These associations are used to protect phase 2 negotiations that are established by using this phase 1. The initiator sends a list of acceptable security associations to the responder in message 1. Each security association defines an acceptable combination of attributes for the ISAKMP SA that is being negotiated. The responder picks a security association that is acceptable and returns the choice to the initiator in message 2.

The following attributes can be negotiated in phase1:

- Authentication method (for example, pre-shared key or RSA signature)
- Hash algorithm (for example, MD5 or SHA1)
- Encryption algorithm (for example, DES, 3DES or AES)
- Diffie-Hellman group information (for example, group 1, group 2, group 5 or group 14)
- Life time and life size of the ISAKMP SA

Messages 3 and 4 are used to exchange information specific to the generation of a shared secret key. This information includes Diffie-Hellman public values and a randomly generated value called a nonce. The initiator sends their Diffie-Hellman public value (for example, g**x mod n) and a nonce in message 3. The responder sends a Diffie-Hellman public value (for example, g**y mod n) and a nonce in message 4. With this information, both the responder and initiator can independently generate the identical keying information. The calculations that are used to generate keying information vary depending on the authentication method that was agreed upon during messages 1 and 2.

The keying information that is generated by both sides includes the following:

- A key that authenticates messages sent under the protection of this ISAKMP SA (for example, phase 2 messages)
- A key that encrypts messages that are sent under the protection of this ISAKMP SA (for example, phase 2 messages)
- Keying material that derives keys that are established for phase 2 SA

Messages 5 and 6 are used to exchange identity information and authentication information. The authentication information varies depending on the authentication method that was agreed upon during messages 1 and 2. For pre-shared key authentication, public key encryption authentication, and revised public key encryption authentication, the information takes the form of an encrypted hash. For signature-based authentication, this information takes the form of a signature. The initiator includes their identity and authentication information in message 5. The responder includes their identity and authentication information in message 6.

Main mode provides a mechanism to exchange certificates when signature-based authentication is used. This mechanism is not shown in Figure 122 on page 972, but works in the following way. In message 5 the initiating ISAKMP server can include the certificate it used to create its signature. In message 6 the responding ISAKMP server might include the certificate it used to create its signature. Inclusion of the certificate is optional unless the ISAKMP server's peer explicitly requests that the certificate be sent.

Aggressive mode

An aggressive mode exchange is comprised of three messages, as shown in Figure 123.

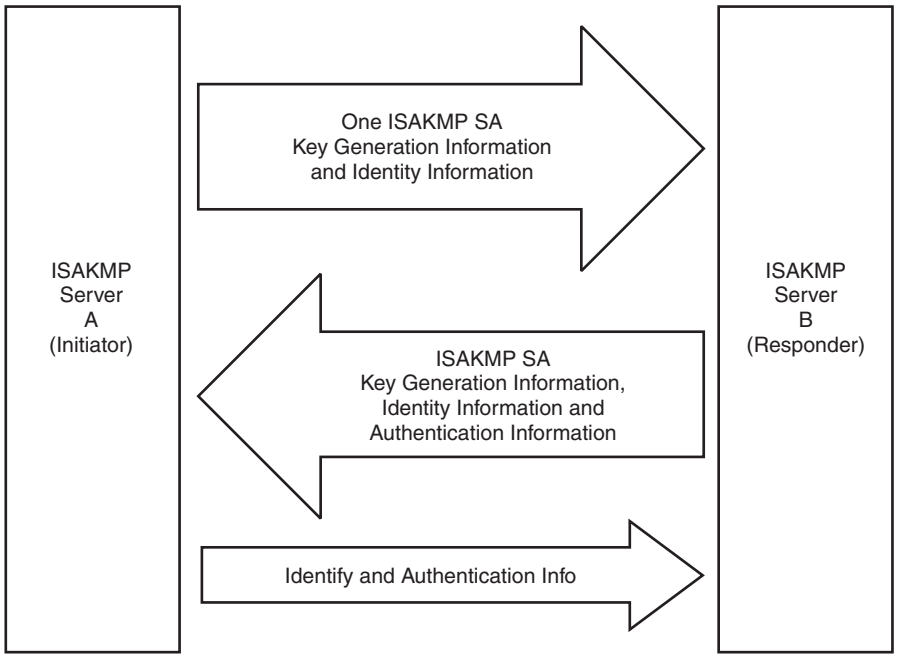

Figure 123. Aggressive mode exchange

Aggressive mode exchanges the same information as Main mode, with the exception of the following:

- In Aggressive mode, the initiator can send only one proposal. In Main mode, the initiator can send a list of proposals.
- In Aggressive mode, only three messages are exchanged instead of six messages as in Main mode.
  - Message 1 of Aggressive mode contains all the information that was contained in messages 1 and 3 of Main mode, plus the identity information sent in message 5 of Main mode.
  - Message 2 of Aggressive mode contains all the information sent in messages 2, 4, and 6 of Main mode.
  - Message 3 of Aggressive mode contains the authentication information that was contained in message 5 of Main mode.
- In Aggressive mode, no messages are required to be encrypted. Message 3 can be sent encrypted, but doing so provides little additional protection. In Main mode, messages 5 and 6 are required to be encrypted. The ISAKMP servers send their identity in messages 5 or 6 of Main mode. The result is that Main mode protects the identity of the ISAKMP servers while Aggressive mode does not. Aggressive mode provides a mechanism to exchange certificates when signature-based authentication is used. This mechanism is not shown in Figure 123 on page 973 but works in the following way. In message 2 the responding ISAKMP server can include the certificate it used to create its signature. In message 3, the initiating ISAKMP server can include the certificate it used to create its signature. Inclusion of the certificates is optional unless the peer of the ISAKMP server explicitly requests that the certificate be sent.

Interpreting IKEv1 daemon phase 1 SA states

The two IKE modes for negotiating phase 1 SAs (main and aggressive) are not themselves negotiable SA attributes. The initiator determines the mode based on the initiator's local policy. The responder can accept or reject the negotiation mode that is selected by the initiator.

Figure 124 on page 975 shows how to interpret phase 1 SA states in Main mode.

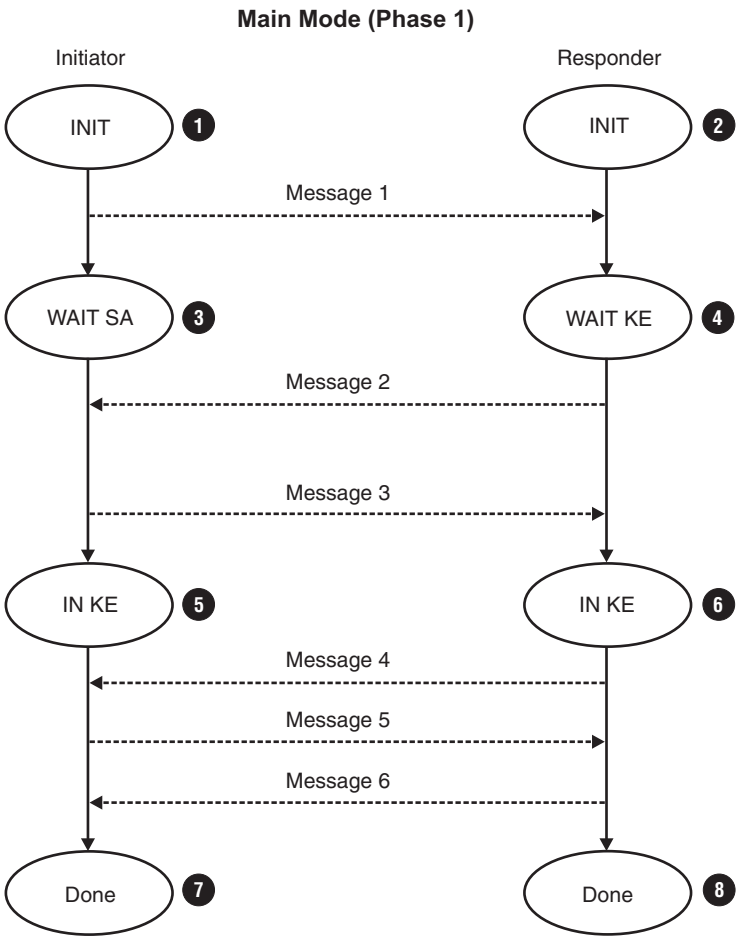

Figure 124. Interpreting phase 1 SA states in Main mode

The following state descriptions apply to the Communications Server IKE daemon when acting as the initiator or responder of a main mode phase 1 SA negotiation (Figure 124). These states are shown in the state field of the **ipsec -k display** command output. See "Main mode" on page 972 for a description of the contents of the messages. The numbers in the following list correspond to the numbered items in Figure 124.

1. The INIT state on the initiator side indicates that message 1 has not yet been sent.
2. The INIT state on the responder side indicates that the responder is processing message 1, which was received from the initiator.
3. This WAIT SA state indicates that the initiator has sent message 1 and is waiting for message 2 from the responder.
4. The WAIT KE state indicates that the responder has processed message 1 and is waiting for message 3 from the initiator.
5. The IN KE state on the initiator side indicates that the initiator has sent message 3.
6. The IN KE state on the responder side indicates that the responder has received message 3.
7. The DONE state on the initiator side indicates that the initiator has received message 6.
8. The DONE state on the responder side indicates that the responder has sent message 6.

Figure 125 shows how to interpret phase 1 SA states in aggressive mode.

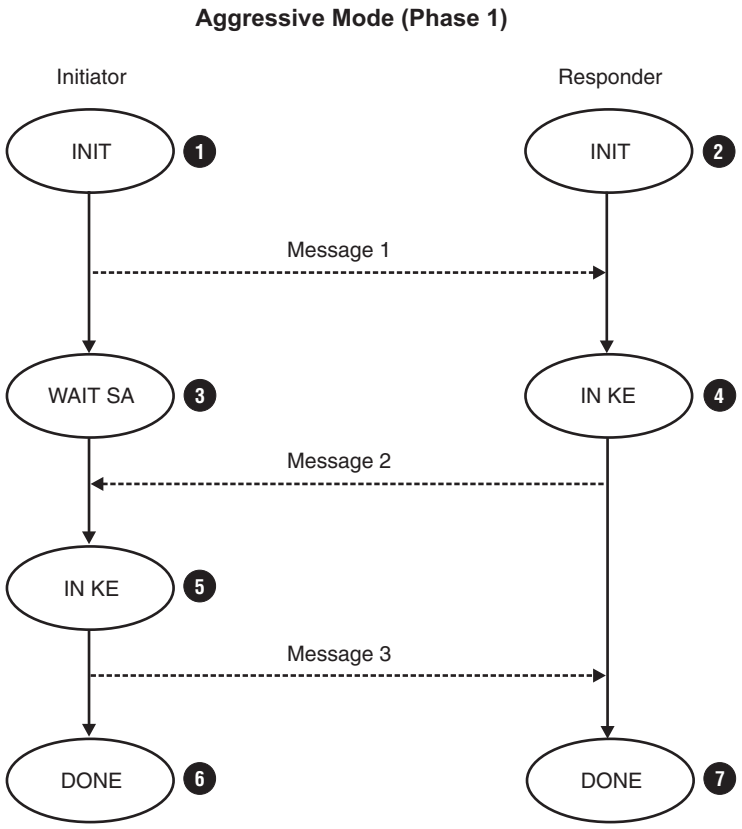

Figure 125. Interpreting phase 1 SA states in Aggressive mode

The following state descriptions apply to the Communications Server IKE daemon when acting as the initiator or responder of an Aggressive mode phase 1 SA negotiation (Figure 125). These states are shown in the state field of the **ipsec -k display** command output. See "Aggressive mode" on page 973 for a description of the contents of the messages. The numbers in the following list correspond to the numbered items inn Figure 125.

1. The INIT state on the initiator side indicates that message 1 has not yet been sent.
2. The INIT state on the responder side indicates that the responder is processing message 1 received from the initiator.
3. The WAIT SA state on the initiator side indicates that the initiator has sent message 1.
4. The IN KE state on the initiator side indicates that the initiator has processed message 1.
5. The IN KE state on the responder side indicates that the responder has received message 2.
6. The DONE state on the initiator side indicates that the initiator has sent message 3.
7. The DONE state on the responder side indicates that the responder has received message 3.

Phase 2

IKE supports one type of phase 2 exchange, Quick mode. Quick mode is an IKE-specific exchange. It is not based on an ISAKMP-defined exchange. Quick mode exchanges are bound to a specific phase1 exchange. This is accomplished by encrypting a hash of each Quick mode message with a cryptographic key derived during the phase 1 exchange. No explicit authentication of the identities involved in a phase 2 exchange is done.

Quick mode

A Quick mode exchange is comprised of three messages, as shown in Figure 126.

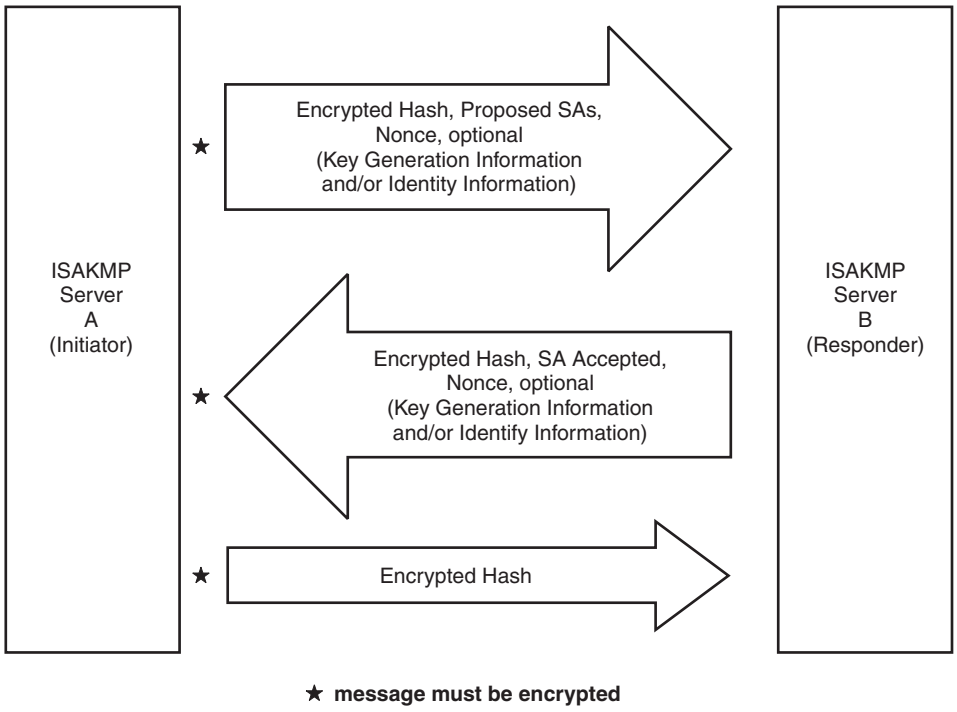

Figure 126. Quick mode exchange messages

In Quick mode, each message contains an encrypted hash. This hash authenticates the source of the message (for example, verifies that it is bound to an ISAKMP SA), authenticates the integrity of the message, and proves liveliness. In message 1, the initiator sends a list of acceptable proposals to the responder. Each proposal defines an acceptable combination of attributes for the non-ISAKMP SA that is being negotiated (AH or ESP SA). The responder picks a proposal that is acceptable and returns the choice to the initiator in message 2.

The attributes that can be negotiated in Quick mode include the following:

- Protocol (AH, ESP, or both AH and ESP)
- Authentication algorithm (for example, Hmac-Md5 or Hmac-Sha)
- Encapsulation mode (tunnel or transport)
- Encryption algorithm (for example, DES, 3DES or AES)
- Diffie-Hellman group information (for example, group 1, group 2, group 5 or group 14)
- Life time and life size of the IPSec SA

Quick mode enables an optional Diffie-Hellman exchange to occur. When the Diffie-Hellman exchange is to take place, the initiator includes a Diffie-Hellman public value (for example, g**x mod n) in message 1, and the responder includes a Diffie-Hellman public value (for example, g**y mod n) in message 2. The key generated from this Diffie-Hellman exchange is used in the calculation that generates the keying material for the non-ISAKMP SA. The Diffie-Hellman exchange provides perfect forward secrecy (PFS).

Quick mode with commit bit

The ISAKMP protocol defines a bit in the ISAKMP message header known as the commit bit. When the commit bit is turned on during a Quick mode exchange, the responder should acknowledge the receipt of message 3. The responder does this by extending the Quick mode exchange to include a fourth message. The major advantage of commit-bit processing is increased interoperability and the elimination of a potential window where IP packets could be dropped during the process of negotiating a new security association.

The z/OS IKE daemon uses commit-bit support as defined in the IKE draft dated May 1999. This draft was written after RFC 2409. No special configuration is required to take advantage of this support. When acting as a responder of a phase 2 negotiation, the IKE daemon always uses commit-bit logic. When acting as an initiator of a phase 2 negotiation, the IKE daemon always honors the commit-bit preference of the responder.

Figure 127 on page 979 shows the new fourth message in the Quick mode exchange, which includes an encrypted hash along with a notify payload indicating that message 3 was received.

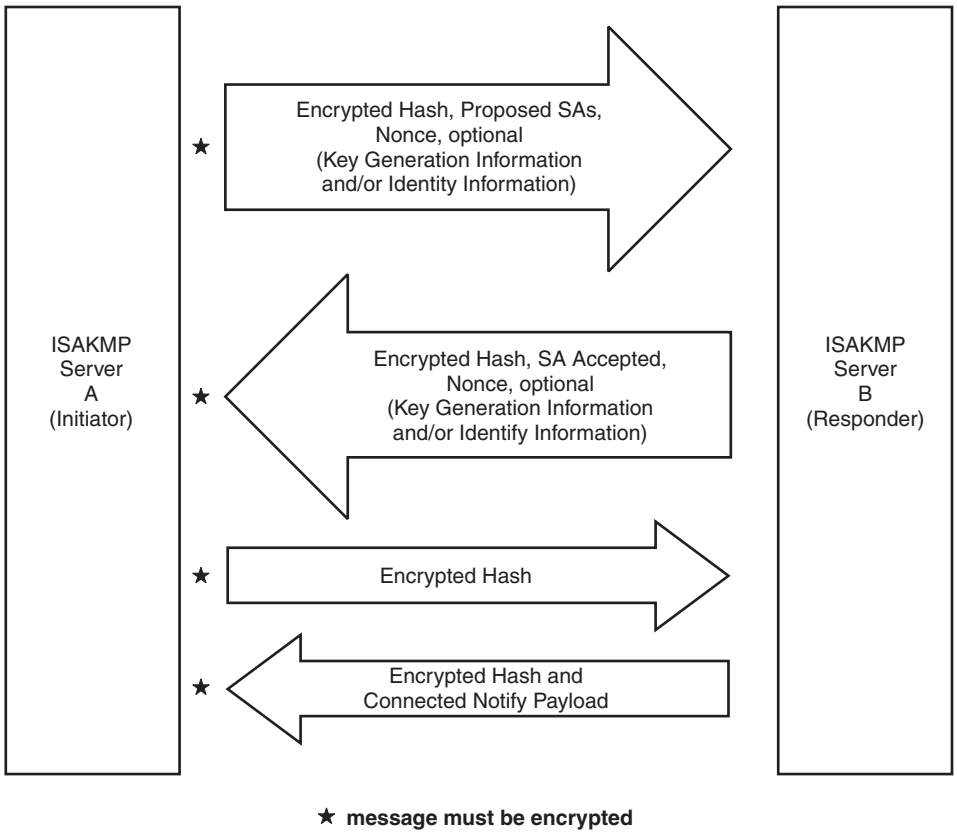

Figure 127. Quick mode exchange with commit-bit support

In a normal Quick mode exchange, the initiator can start using a newly negotiated SA immediately after sending message 3. The responder does not start using the newly negotiated SA until it receives message 3. Message 3 is sent using UDP. Because UDP is not a reliable protocol, it is possible that the initiator sends message 3 and that this message never gets processed by the responder. In this case, the responder retransmits message 2 back to the initiator, causing the initiator to retransmit message 3. Unfortunately, during the time between such retransmissions, the initiator might start using the SA to protect an IP packet. Any such packet would be discarded by the responder until it successfully processed message 3.

In a Quick mode exchange with commit processing, the initiator defers the usage of a newly negotiated SA until one of the following events occur:

- The initiator receives a connected notify message
- The initiator receives an IP packet that was protected with the SA

The responder continues to start using the newly negotiated SA when it receives message 3. This eliminates the window where one side might start using an SA before the other side knows that it is safe to use the SA.

On z/OS, an SA is considered to be in a pending state while the initiator is waiting for a connected notify message (for example, message 4). An SA is placed into a pending state only if another SA that could be used to protect outbound traffic exists. An SA in pending state remains in pending state until one of the following events occur:

- A connected notify is received

- A message protected by the SA is received
- The last usable SA expires

Interpreting IKEv1 daemon phase 2 SA states

Commit-bit support is not a negotiable phase 2 SA attribute. The Communications Server IKE daemon always includes the commit bit when initiating a Quick mode negotiation. If the responder does not support commit-bit processing, the Communications Server IKE daemon does not wait for a connected notify message from the responder. If the initiator does not have commit-bit support, then the Communications Server IKE daemon does not send a connected notify message when acting as the responder.

Quick mode (phase 2) SA states without commit-bit support:
Figure 128 shows interpreting Quick mode (phase 2) SA states without commit-bit support

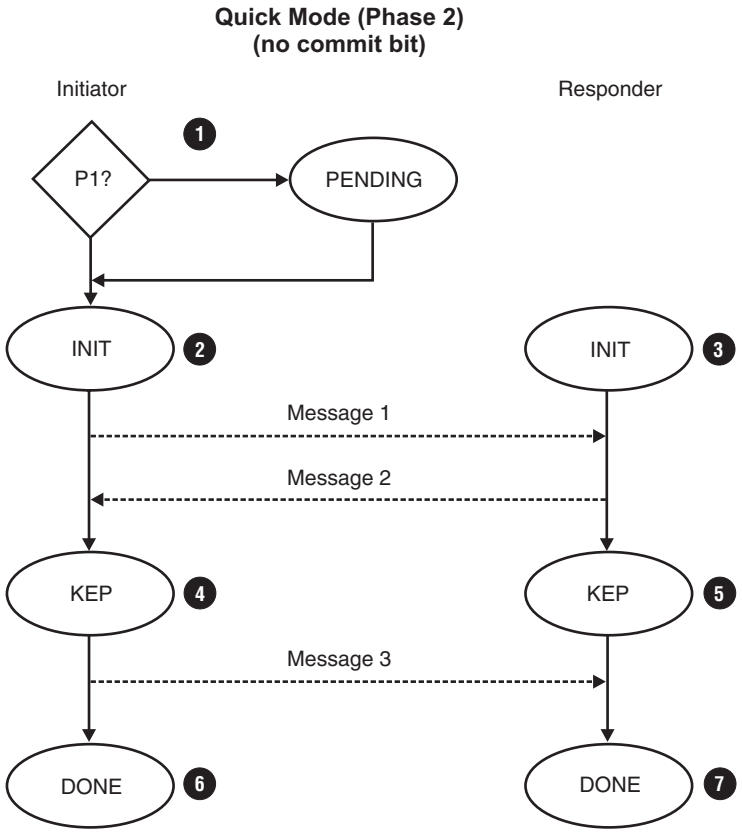

Figure 128. Quick (phase 2) SA states without commit-bit support

The following state descriptions apply to the Communications Server IKE daemon when acting as the initiator or responder of a Quick mode (phase 2) SA negotiation without commit-bit support (Figure 128). These states are shown in the state field of the **ipsec -y display -b** command output. See "Quick mode" on page 977 for a description of the contents of the messages.

1. The INIT state on the initiator side indicates that message 2 has not yet been received.
2. The INIT state on the responder side indicates that the responder has not yet sent message 2.

3. A phase 1 SA must be established between the initiator and responder before the initiator can send message 1 of Quick mode. The PENDING state indicates that the initiator is waiting for a phase 1 negotiation to complete with the responder. For more information, "Interpreting IKEv1 daemon phase 1 SA states" on page 974. After the phase 1 negotiation completes, message 1 of Quick mode can be sent.
4. The KEP state on the initiator side indicates that message 2 has been received.
5. The KEP state on the initiator side indicates that message 2 has been sent.
6. The DONE state on the initiator side indicates that message 3 has been sent.
7. The DONE state on the responder side indicates that message 3 has been received.

Quick mode (phase 2) SA states with commit-bit support:
Figure 129 shows interpreting Quick mode (phase 2) SA states with commit-bit support.

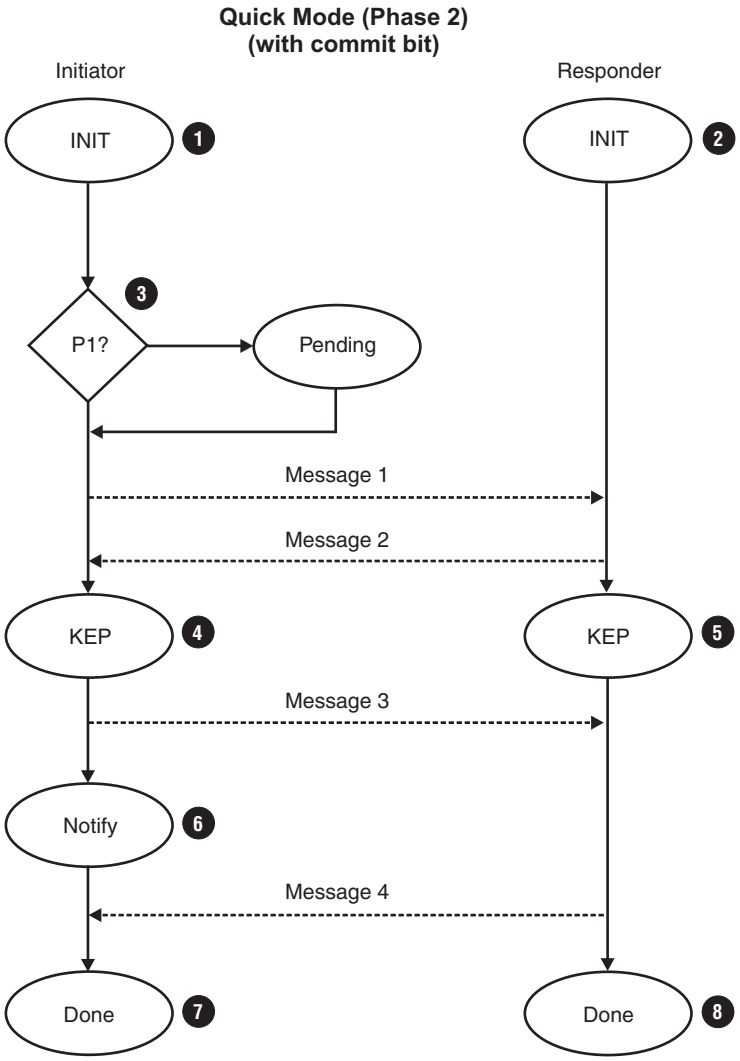

Figure 129. Quick (phase 2) SA states with commit-bit support

The following state descriptions apply to the Communications Server IKE daemon when acting as the initiator or responder of a Quick mode (phase 2) SA negotiation

with commit-bit support (Figure 129 on page 981). These states are shown in the state field of the **ipsec -y display -b** command output. See "Quick mode with commit bit" on page 978 for a description of the contents of the messages.

1. The INIT state on the initiator side indicates that message 2 has not yet been received.
2. The INIT state on the responder side indicates that the responder has not yet sent message 2.
3. A phase 1 SA must be established between the initiator and responder before the initiator can send message 1 of Quick mode. The PENDING state indicates that the initiator is waiting for a phase 1 negotiation to complete with the responder. For more information, see "Interpreting IKEv1 daemon phase 1 SA states" on page 974. After the phase 1 negotiation completes, message 1 of Quick mode can be sent.
4. The KEP state on the initiator side indicates that message 2 has been received.
5. The KEP state on the responder side indicates that message 2 has been sent.
6. The NOTIFY state indicates that the initiator has sent message 3 and is waiting for message 4.
7. The DONE state on the initiator side indicates that message 4 has been received.
8. The DONE state on the responder side indicates that message 4 has been sent.

ISAKMP Main mode limitations

This topic contains information about three Main mode scenarios.

Main mode scenario 1

Key policy definition is based on the identities of remote ISAKMP servers. Unfortunately, during a Main mode exchange the responding ISAKMP server must accept a key proposal prior to learning the identity of the initiating ISAKMP server. The responder must later verify that the proposal that was agreed to is acceptable with defined policy when the identity becomes known.

The z/OS IKE daemon handles this limitation as follows:

1. Upon receipt of message 1, the IKE daemon uses the IP address of the initiator and responder to find an applicable KeyExchangeRule, which encapsulates the key policy:
   - If an applicable KeyExchangeRule is found, it is considered tentative until the identity of the initiator becomes known.
2. Upon receipt of message 5, which includes the initiator's identity, the IKE daemon uses the IP address of the initiator, the IP address of the responder, and the identity of the initiator to find an applicable KeyExchangeRule:
   - If a KeyExchangeRule is not found or is found but is inconsistent with the proposal accepted in message 1, the negotiation fails.
   - If a KeyExchangeRule is found and is consistent with the proposal accepted in message 1, it is considered final, and the negotiation proceeds.

Main mode scenario 2

Pre-shared keys are defined based on the identities of ISAKMP servers. Ideally, pre-shared keys are unique between ISAKMP server pairs. Unfortunately, during a Main mode exchange the responding ISAKMP server must determine the pre-shared key to use before learning the identity of the initiating ISAKMP server.

The z/OS IKE daemon handles this limitation as follows:

1. A key proposal is selected as described in "Main mode scenario 1" on page 982.
2. If the selected key proposal indicates pre-shared key mode authentication, then the IKE daemon must use a pre-shared key to generate message 4.
3. Upon receipt of message 5, the IKE daemon must use the same pre-shared key to decrypt the message to learn the identity of the initiating ISAKMP server.
4. After message 5 is successfully decrypted, the IKE daemon uses the IP address of the initiator, the IP address of the responder, and the identity of the initiator to find an applicable KeyExchangeRule:
   - If a KeyExchangeRule is not found or is found but is inconsistent with the proposal accepted in message 1, the negotiation fails.
   - If a KeyExchangeRule is found and is consistent with the proposal accepted in message 1, it is considered final, and the negotiation proceeds.

Main mode scenario 3

Certificate authority (CA) certificates are associated with the identities of remote ISAKMP servers. When RSA signature mode authentication is being performed, the ISAKMP responder might send one or more certificate requests to the ISAKMP initiator to guide the initiator in selecting a certificate signed by an acceptable CA. Unfortunately, during a Main mode exchange the responding ISAKMP server must send a certificate request before learning the identity of the initiating ISAKMP server.

The z/OS IKE daemon handles this limitation as follows:

1. A key proposal is selected as described in Scenario 1.
2. If the selected key proposal indicates RSA signature mode authentication, then the IKE daemon includes one or more certificate requests in message 4.
   - If a tentative KeyExchangeRule is in effect and the KeyExchangeRule's RemoteSecurityEndpoint includes one or more CaLabels, a certificate request corresponding to each CaLabel is included in message 4.
   - If the RemoteSecurityEndpoint does not include a CaLabel, a certificate request corresponding to each SupportedCertAuth is included in message 4.
   - If there are no applicable CaLabels or SupportedCertAuth statements configured, an empty certificate request is included in message 4, indicating that the initiator can use a certificate signed by any CA.

IKE version 2 protocol

IKE version 2 (IKEv2) is defined in RFC 5996, and is intended to replace IKEv1. This topic describes how the IKEv2 protocols are used to negotiate security associations (SAs) and exchange keys between two systems that want to communicate securely.

Overview of negotiating IKEv2 security associations

The IKEv2 protocol is very similar to IKEv1 in many respects. Both protocols establish SAs in two phases. They first establish an SA that securely carries IKE messages between the peers, and subsequently establish additional SAs to carry the protected ESP or AH traffic. For IKEv2, the SA that carries IKE messages is referred to as the IKE SA, and the SAs for ESP and AH are child SAs.

For IKEv1, the corresponding terms for the two types of SAs are "ISAKMP SA" and "IPSec SA". We use the terms "phase 1 SA" and "phase 2 SA" to refer to the two SA types when the version of IKE is unknown or unimportant.

The message formats defined for IKEv2 are very similar to those for IKEv1. Both formats start with a message header that contains a protocol version field, so a receiving node can receive both types of messages on a single UDP port (by default, port 500), and easily tell whether the message is IKEv1 or IKEv2. Both include variable length payloads of a similar generic format.

The specific content and sequences of messages for IKEv2 are quite different from IKEv1. In IKEv2, all communications consist of pairs of messages: a request and a response. The pair is called an "exchange". The initial exchanges consist of the IKE_SA_INIT exchange and the IKE_AUTH exchange. These two exchanges establish both the IKE SA and the first Child SA. Subsequent exchanges are the CREATE_CHILD_SA exchanges and INFORMATIONAL exchanges, which perform other duties such as establishing additional child SAs and deleting SAs.

Initial exchanges

Activation of an IKE_SA requires completion of two exchanges, IKE_SA_INIT exchange and IKE_AUTH exchange, as illustrated in Figure 130.

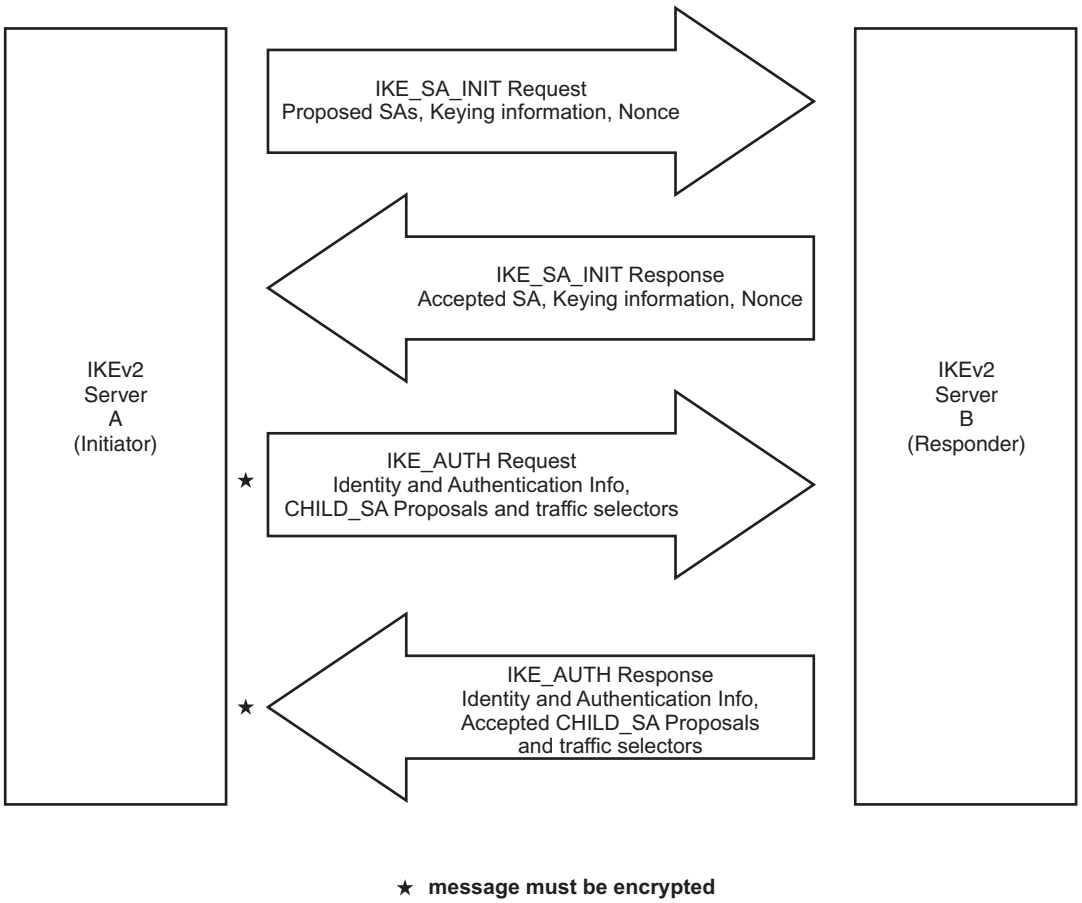

Figure 130. IKEv2 initial exchanges

The first exchange of an IKEv2 activation attempt is the IKE_SA_INIT exchange. The initiator sends a list of security association proposals to the responder in the IKE_SA_INIT request. Each proposal defines a combination of attributes for the IKE SA that is being negotiated. The initiator also includes its nonce and Diffie-Hellman value in the IKE_SA_INIT request. The responder picks a proposal

that is acceptable and returns its choice to the initiator in the IKE_SA_INIT response, along with its own nonce and Diffie-Hellman value.

The following attributes of the IKE SA can be negotiated during the IKE_SA_INIT exchange:

- Message authentication algorithm (for example, HMAC-MD5-96 or HMAC-SHA1-96)
- Pseudo-random function for key generation (for example, HMAC-MD5 or HMAC-SHA1)
- Encryption algorithm (for example, DES, 3DES or AES)
- Diffie-Hellman group information (for example, group 1, group 2, group 5 or group 14)

Once the IKE_SA_INIT exchange completes successfully, both the responder and initiator can independently generate the identical keying information that supports the IKE SA. This keying information includes the following:

- A pair of keys used to authenticate the IKE peers
- A pair of keys used to authenticate messages sent under the protection of this IKE SA
- A pair of keys used to encrypt messages that are sent under the protection of this IKE SA
- Keying material that derives keys that are established for child SAs

Note that, unlike IKEv1, each IKEv2 peer chooses its own authentication method. On a single IKE SA, one peer might choose to be authenticated using a pre-shared key, while the other peer chooses digital signature authentication. If the two peers both choose pre-shared key as the authentication method, the IKEv2 protocol allows their keys to be different, but the z/OS implementation requires both peers to use the same key. Also, the IKE SA life time and life size are not negotiated between the two IKEv2 peers. Each peer manages its own independent value of life time and life size for each IKE SA.

In some cases, negotiation of these attributes may require more than one IKE_SA_INIT exchange. The initiator makes a guess as to which proposal the responder will choose, and sends a Diffie-Hellman value that corresponds to that guess. If the responder chooses a different proposal, it rejects the IKE_SA_INIT, and includes an indication in the response that identifies its chosen proposal. The initiator learns the correct choice from the IKE_SA_INIT response, and initiates a new IKE_SA_INIT exchange with the responder's chosen proposal and corresponding Diffie-Hellman value.

IKEv2 also provides a mechanism to exchange certificates when signature-based authentication is used. In the IKE_AUTH request the initiator can include the certificate it used to create its signature. In the IKE_AUTH response, the responder can include the certificate it used to create its signature. Inclusion of the certificates is optional. IKEv2 can be configured to send and/or receive new certificate encoding types that are not supported for IKEv1:

- Hash and URL of an X.509 certificate
- Hash and URL of an X.509 bundle

Because both peers have all the required keying information, all but the headers of all subsequent requests and responses sent on the IKE SA, by either peer, are encrypted and authenticated.

The second exchange is the IKE_AUTH exchange. This exchange completes the activation of the IKE SA, and also sets up an SA for the first (and often only) AH or ESP child SA. Details of first child SA activation are described in “First Child SA” on page 988.

Note: The z/OS IKE daemon uses the Network Security Services (NSS) Certificate service for creating and verifying digital signatures during the IKE_AUTH exchange.

To complete activation of the IKE SA, the initiator transmits an IKE_AUTH request that contains its identity and authentication information. The authentication information varies depending on the initiator's authentication method that was declared in the IKE_SA_INIT request. For pre-shared key authentication, the information takes the form of an encrypted hash. For signature-based authentication, this information takes the form of a digital signature. The initiator may also include the expected identity of the responder in the IKE_AUTH request. This is useful when the machine on which the responder is running is hosting multiple identities at the same IP address. The responder includes its identity and authentication information in the IKE_AUTH response.

When a certificate payload of one of these encoding types is received, the actual certificate is retrieved from an HTTP server using the provided URL, and the retrieved certificate is verified using the provided hash. The hash and URL are typically smaller than the certificate they refer to, so some efficiency is gained by using them. However, there is additional cost to retrieve the certificate from the HTTP server.

Interpreting IKEv2 IKE SA states

Figure 131 on page 987 shows how to interpret IKE SA states.

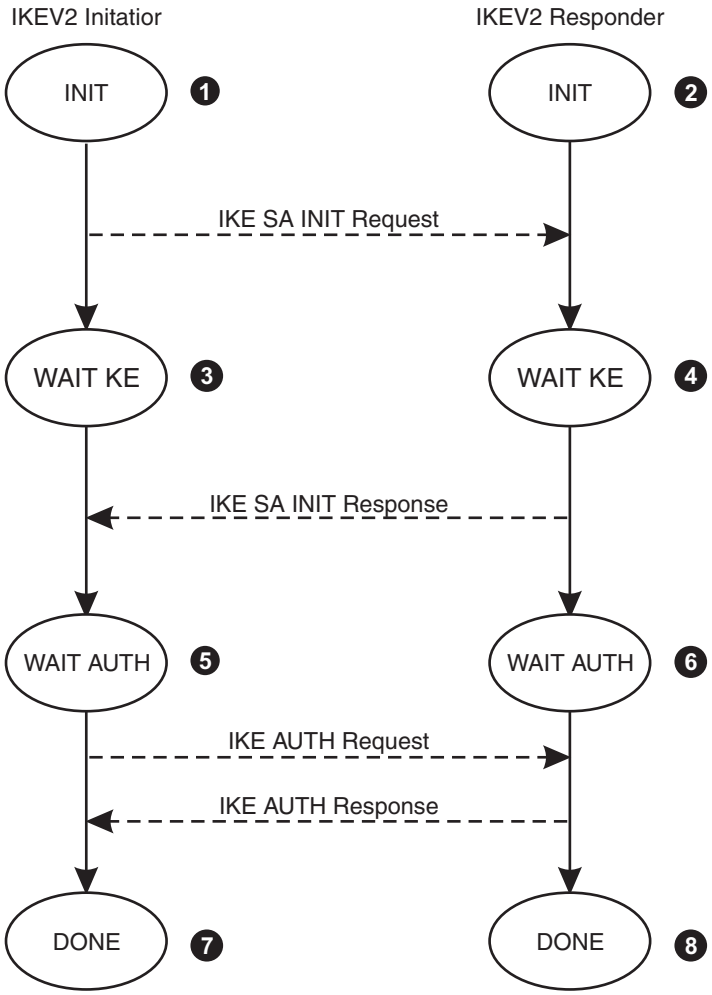

Figure 131. Interpreting IKEv2 IKE SA states

The following state descriptions apply to the Communications Server IKE daemon when acting as the initiator or responder of an IKEv2 phase 1 SA negotiation. These states are shown in the state field of the **ipsec -k display** command output. See "Initial exchanges" on page 984 for a description of the contents of the messages. The numbers in the following list correspond to the numbered items in Figure 131.

1. The INIT state on the initiator side indicates that the IKE_SA_INIT request has not yet been sent.
2. The INIT state on the responder side indicates that the responder is processing the IKE_SA_INIT request, which was received from the initiator.
3. This WAIT KE state indicates that the initiator has sent the IKE_SA_INIT request and is waiting for the IKE_SA_INIT response from the responder.
4. The WAIT KE state indicates that the responder has processed the IKE_SA_INIT and is waiting for the IKE_AUTH request from the initiator.
5. The WAIT AUTH state on the initiator side indicates that the initiator has sent the IKE_AUTH request
6. The WAIT AUTH state on the responder side indicates that the responder has received the IKE_AUTH request
7. The DONE state on the initiator side indicates that the initiator has received the IKE_AUTH response

8. The DONE state on the responder side indicates that the responder has sent the IKE_AUTH response

Child SA activation

For IKEv2, activation of the first Child SA under an IKE_SA is handled slightly differently than activation of subsequent Child SAs under that same IKE SA.

First Child SA

The IKEv2 protocol was designed so that the first Child SA is activated during processing of the IKE_AUTH request and response. For many configurations, this means that the IKE SA and Child SA are both activated by only four messages.

The IKE_AUTH request contains the initiator's list of SA proposals, and the traffic selectors that describe the traffic to be protected by the Child SA. However, the IKE_AUTH request does NOT contain keying information or a nonce that is specific to the Child SA. The nonces and keying information from the IKE_SA_INIT exchange are used in computing the keys for the first Child SA. See Figure 130 on page 984 for an illustration of the IKE_AUTH exchange.

Processing of the SA proposals and the traffic selectors during the IKE_AUTH exchange is the same as in CREATE_CHILD_SA processing, described in "Additional Child SAs."

Additional Child SAs

Each additional Child SA is established using a single CREATE_CHILD_SA exchange, as illustrated in Figure 132.

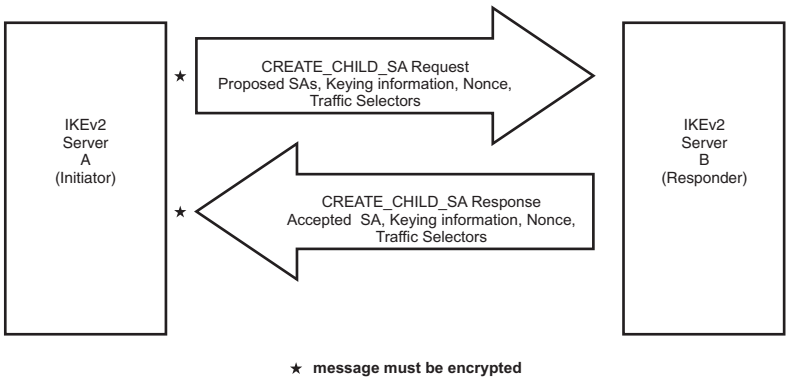

Figure 132. IKEv2 CREATE_CHILD_SA exchange

The initiator sends a CREATE_CHILD_SA request, containing a list of acceptable proposals for the Child SA. Each proposal defines an acceptable combination of attributes for the Child SA that is being negotiated (AH or ESP SA). The responder picks a proposal that is acceptable and returns the choice to the initiator in the CREATE_CHILD_SA response. The attributes that can be negotiated include the following:

- Protocol (AH or ESP)
- Authentication algorithm (for example, HMAC-MD5 or HMAC-SHA)
- Encapsulation mode (tunnel or transport)
- Encryption algorithm (for example, DES, 3DES or AES)
- Diffie-Hellman group information (for example, group 1, group 2, group 5 or group 14)

Unlike IKEv1, the Child SA life time and life size are not negotiated between the two IKEv2 peers. Each peer manages its own independent value of life time and life size for each Child SA.

An optional Diffie-Hellman exchange may occur during the CREATE_CHILD_SA exchange. When the Diffie-Hellman exchange is to take place, the initiator includes a Diffie-Hellman public value in the CREATE_CHILD_SA request, and the responder includes a Diffie-Hellman public value in the CREATE_CHILD_SA response. The key generated from this Diffie-Hellman exchange is used in the calculation that generates the keying material for the Child SA. The Diffie-Hellman exchange provides perfect forward secrecy (PFS), which ensures the Child SA keys are derived independently from the IKE SA keys.

Traffic selectors that describe the traffic to be protected by the SA are also negotiated during the CREATE_CHILD_SA exchange. The initiator sends a set of proposed traffic selectors in the CREATE_CHILD_SA request, and the responder can narrow the traffic selection by sending a subset of the initiator's proposed traffic selectors on the CREATE_CHILD_SA response.

As with IKE_SA_INIT, in some cases, negotiation of these attributes may require more than one CREATE_CHILD_SA exchange. The initiator makes a guess as to which proposal the responder will choose, and sends a Diffie-Hellman value that corresponds to that guess. If the responder chooses a different proposal, it rejects the CREATE_CHILD_SA, and includes an indication in the response that identifies its chosen proposal. The initiator learns the correct choice from the CREATE_CHILD_SA response, and initiates a new CREATE_CHILD_SA exchange with the responder's chosen proposal and corresponding Diffie-Hellman value

Interpreting IKEv2 Child SA states

Figure 133 on page 990 shows how to interpret IKE_SA states.

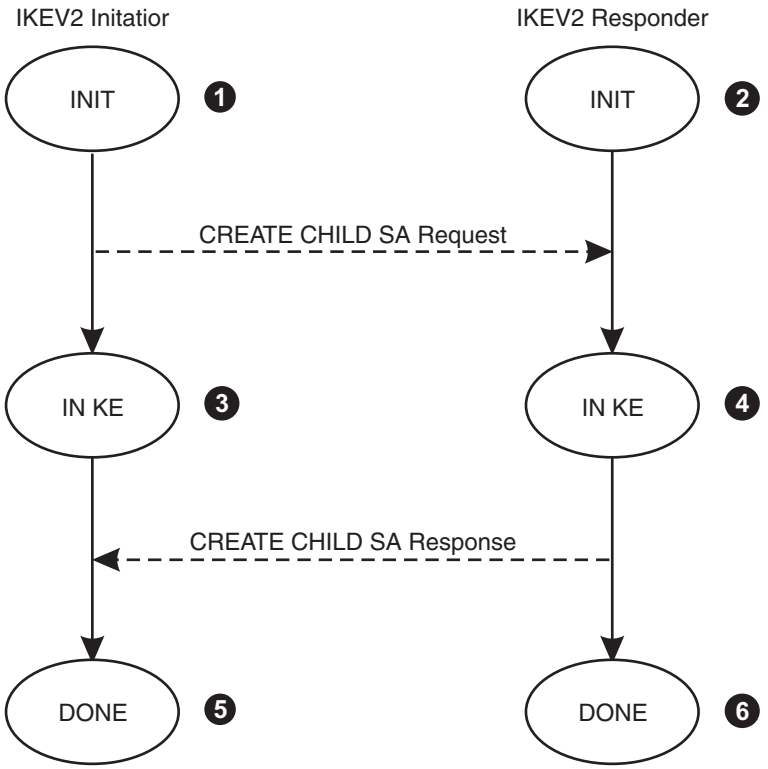

Figure 133. IKEv2 Child SA states

The following state descriptions apply to the Communications Server IKE daemon when acting as the initiator or responder of an IKEv2 phase 2 SA negotiation. These states are shown in the state field of the **ipsec -y display -b** command output. See "Child SA activation" on page 988 for a description of the contents of the messages. The numbers in the following list correspond to the numbered items in Figure 133.

1. The INIT state on the initiator side indicates that the CREATE_CHILD_SA Request has not yet been sent.
2. The INIT state on the responder side indicates that the responder is processing the CREATE_CHILD_SA Request, which was received from the initiator.
3. This IN KE state indicates that the initiator has sent the CREATE_CHILD_SA Request and is waiting for the CREATE_CHILD_SA Response from the responder.
4. The IN KE state indicates that the responder is processing the CREATE_CHILD_SA
5. The DONE state on the initiator side indicates that the initiator has received the CREATE_CHILD_SA Response
6. The DONE state on the responder side indicates that the responder has sent the CREATE_CHILD_SA Response

Key exchange limitations

Key policy definition is based on the identities of remote IKE servers. During an IKEv2 authentication exchange, the IKE daemon must choose a key exchange rule and action before the identity of the remote IKE server is known. This determines policy decisions that are made early in the negotiation. These decisions cannot be reversed later in the negotiation; however, where applicable, once the identity of

the remote IKE server is known these policy decisions are verified. The z/OS IKE daemon manages the policy selection process as described in the following topics:

- "Responder"
- "Initiator"

Responder

When z/OS IKED is acting as responder:

- Upon receipt of the IKE_SA_INIT request message, the IKE daemon uses the IP addresses of the initiator and responder to find the first applicable KeyExchangeRule that encapsulates these addresses. If a rule is found, it is considered to be tentative until the identity of the initiator becomes known. All KeyExchangePolicy settings applicable to the IKE_SA_INIT and IKE_AUTH exchanges are determined from this tentative rule, including SA proposal attributes, pre-shared key, and certificate authority (CA) labels.
- Upon receipt of the IKE_AUTH request message, which includes the initiator's identity (IDi):
  - If the request message also contains the optional requested responder identity (IDr), the IKE daemon uses the IP addresses of the initiator and responder, plus both the initiator and responder identities contained in the request message, to find an applicable KeyExchangeRule. At this point:
    - If a KeyExchangeRule is not found, the negotiation continues to step B below.
    - If a KeyExchangeRule is found, and the SA is using pre-shared key identity protection in either direction, and the KeyExchangeRule's pre-shared key does not match the pre-shared key of the tentative KeyExchangeRule, the negotiation continues to step B below.
    - If a KeyExchangeRule is found, but it does not include an SA proposal that is consistent with the proposal accepted in the IKE_SA_INIT exchange, the negotiation continues to step B below.
    - Otherwise, the new KeyExchangeRule is considered final, and the negotiation proceeds. This KeyExchangeRule may contain a local security endpoint identity that is different from the tentative KeyExchangeRule.
  - If the request message does not contain the optional requested responder identity (IDr), or if it does but step A above fails, the IKE daemon uses the IP addresses of the initiator and responder, plus the identity of the initiator, to find an applicable KeyExchangeRule. At this point:
    - If a KeyExchangeRule is not found, the negotiation fails.
    - If a KeyExchangeRule is found, and the SA is using pre-shared key identity protection in either direction, and the KeyExchangeRule's pre-shared key does not match the pre-shared key of the tentative KeyExchangeRule, the negotiation fails.
    - If a KeyExchangeRule is found, but it does not include an SA proposal that is consistent with the proposal accepted in the IKE_SA_INIT exchange, the negotiation fails.
    - Otherwise, the new KeyExchangeRule is considered final, and the negotiation proceeds. This KeyExchangeRule may contain a local security endpoint identity that is different from the tentative KeyExchangeRule.

Initiator

When z/OS IKED is acting as initiator:

- Prior to sending the IKE_SA_INIT request, the initial KeyExchangeRule is selected according to the rules specified in the IpLocalStartAction statement. See

z/OS Communications Server: IP Configuration Reference for more information. If a rule is found, it is considered to be tentative until the identity of the responder becomes known. All KeyExchangePolicy settings applicable to the IKE_SA_INIT and IKE_AUTH exchanges are determined from this tentative rule, including SA proposal attributes, pre-shared key, and certificate authority (CA) labels.

- The initiator identity (IDi) that the IKE daemon sends in the IKE_AUTH request message is determined from this, and if applicable, the optional responder identity (IDr) sent in the IKE_AUTH request is also determined from this.
- Upon receipt of the IKE_AUTH response message, which includes the responder's identity (IDr), the IKE daemon uses the IP addresses of the initiator and responder, the initiator identity determined from the tentative KeyExchangeRule, plus the responder identity contained in the response message, to find an applicable KeyExchangeRule. At this point:
  - If a KeyExchangeRule is not found, the negotiation fails.
  - If a KeyExchangeRule is found, and the SA is using pre-shared key identity protection in either direction, and the KeyExchangeRule's pre-shared key does not match the pre-shared key of the tentative KeyExchangeRule, the negotiation fails.
  - If the tentative KeyExchangeRule was located using a RemoteSecurityEndpoint configured on an IpLocalStartAction statement, and if the remote IKE daemon's identity (IDr) is not encompassed by the tentative RemoteSecurityEndpoint identity, the negotiation fails.
  - If a KeyExchangeRule is found, but it does not include an SA proposal that is consistent with the proposal accepted in the IKE_SA_INIT exchange, the negotiation fails.
  - Otherwise, the new KeyExchangeRule is considered final, and the negotiation proceeds.

Traversing a NAT

There are several incompatibility issues that exist between IPSec and Network Address Translation (NAT). These incompatibility issues are described in RFC 3715, *IPsec-Network Address Translation (NAT) Compatibility Requirements*. Three RFCs were written to address these incompatibility issues:

- RFC 3947, *Negotiation of NAT-Traversal in the IKE*
- RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
- RFC 5996, *Internet Key Exchange (IKEv2) Protocol*

All of these RFCs have been implemented on z/OS, providing z/OS with the capability to perform IPSec while traversing a NAT in a limited set of environments. RFC 3947 augments IKE's Main mode, Aggressive mode, and Quick mode messages flows to include additional information. It also provides for the negotiation of two new encapsulation modes.

To provide the possibility of interoperability with some pre-RFC 3947 implementations z/OS also provides support for the following pre-RFC 3947, *Negotiation of NAT-Traversal in the IKE* drafts:

- draft-ietf-ipsec-nat-t-ike-02
- draft-ietf-ipsec-nat-t-ike-03

Impacts to IKEv1 phase 1 (Main and Aggressive mode)
RFC 3947 requires that a vendor ID payload containing a NAT traversal

vendor ID be exchanged between two IKEv1 peers. The vendor ID payload is an existing ISAKMP payload. The vendor ID payload is used by an IKE daemon to advertise support for a feature that is an extension to RFC 2408 (ISAKMP) and RFC 2409 (IKE). The vendor ID that is contained in the payload identifies the feature. The NAT traversal vendor ID is defined to be an MD5 hash of the vendor string RFC 3947.

The NAT traversal vendor ID must be received before an IKE daemon can send any of the new payloads and encapsulation modes that are defined in RFC 3947. Likewise, an IKE daemon should not send any of the new payloads and encapsulation modes defined in RFC 3947 without first sending the NAT traversal vendor ID.

If the initiator of an IKEv1 phase 1 negotiation wants to advertise support for RFC 3947, it must send the NAT traversal vendor ID in message 1 of a Main mode exchange or message 1 of an Aggressive mode exchange. If the responder of an IKEv1 phase 1 negotiation wants to advertise support for RFC 3947, it must send the NAT traversal vendor ID in message 2 of a Main mode exchange or message 2 of an Aggressive mode exchange.

z/OS provides limited support for several pre-RFC 3947 drafts, as well as additional z/OS-to-z/OS NAT traversal capabilities. Unique vendor IDs are used to identify these various levels of NAT traversal support. Table 108 shows the NAT traversal vendor IDs that are recognized by z/OS. The vendor IDs are listed from least functional to most functional. If z/OS receives multiples of these IDs, it uses the most functional level of support that it received.

Table 108 lists vendor ID strings.

Table 108. Vendor ID

Vendor ID string	Vendor ID
draft-ietf-ipsec-nat-t-ike-02\n	90cb8091 3ebb696e 086381b5 ec427b1f
draft-ietf-ipsec-nat-t-ike-02	cd604643 35df21f8 7cfdb2fc 68b6a448
draft-ietf-ipsec-nat-t-ike-03	7d9419a6 5310ca6f 2c179d92 15529d56
RFC 3947	4a131c81070358455c5728f20e95452f
z/OS CS-IKE NAT Traversal Level 1	95305bb5 64b82a30b 66968bbc 5326a8d

In z/OS, NAT traversal support can be enabled or disabled with the AllowNat parameter. The AllowNat parameter can be specified on the KeyExchangePolicy statement, the KeyExchangeAction statement of the IPSec Policy file, or both. When AllowNat is set to `NO` the z/OS IKE daemon does not send NAT traversal vendor IDs. See z/OS Communications Server: IP Configuration Reference for additional details about the AllowNat parameter.

RFC 3947 defines a mechanism for discovering the existence of NAT devices residing between two IKE daemons using IKEv1, as well as the location of the NAT devices. This mechanism is the NAT Discovery (NAT-D) payload. The NAT-D payload is an extension to RFC 2408 and 2409. It contains a hash of several pieces of information including an IP address and port value from the IP packet that is being sent to an IKE peer (for example, the packet containing the NAT-D payload).

Each IKEv1 peer sends two or more NAT-D payloads. The destination IP address and port of the outbound IKE packet are used to construct the hash that is contained within the first NAT-D payload. The source IP

address and port of the outbound IKE packet are used to construct the hash that is contained within the second NAT-D payload. Normally, only two NAT-D payloads are exchanged; however, if the sender of the packet has multiple IP addresses and it does not know which IP address is used to send the packet, it can send a NAT-D payload for each IP address it owns.

The initiator of an IKEv1 phase 1 negotiation must send its NAT-D payloads in message 3 of a Main mode exchange or message 3 of an Aggressive mode exchange. The responder of an IKEv1 phase 1 negotiation must send its NAT-D payloads in message 4 of a Main mode exchange or message 2 of an Aggressive mode exchange.

Impacts to IKEv2 phase 1 (IKE_SA_INIT)

RFC 5996 requires that an IKE daemon that supports IKEv2 NAT traversal must send NAT_DETECTION_SOURCE_IP and NAT_DETECTION_DESTINATION_IP payloads.

In z/OS, you can enable or disable NAT traversal support with the AllowNat parameter. You can specify the AllowNat parameter on the KeyExchangePolicy statement, the KeyExchangeAction statement of the IPSec Policy file, or both. When AllowNat is set to NO, the z/OS IKE daemon does not send NAT_DETECTION_SOURCE_IP or NAT_DETECTION_DESTINATION_IP payloads. See z/OS Communications Server: IP Configuration Reference for additional details about the AllowNat parameter.

If the initiator of an IKEv2 phase 1 negotiation wants to advertise support for IKEv2 NAT Traversal, it must send NAT_DETECTION_SOURCE_IP and NAT_DETECTION_DESTINATION_IP payloads in the IKE_SA_INIT request. If the responder of an IKEv2 phase 1 negotiation wants to advertise support for IKEv2 NAT Traversal, it must send NAT_DETECTION_SOURCE_IP and NAT_DETECTION_DESTINATION_IP payloads in the IKE_SA_INIT response.

RFC 5996 defines a mechanism for discovering the existence of NAT devices residing between two IKE daemons utilizing IKEv2, as well as the location of the NAT devices. This mechanism is based upon the NAT_DETECTION_SOURCE_IP and NAT_DETECTION_DESTINATION_IP payloads. The NAT_DETECTION_SOURCE_IP payload contains a hash of several pieces of information including the source IP address and source port value from the IP packet that is being sent to an IKE peer (for example, the packet containing the NAT_DETECTION_SOURCE_IP payload). The NAT_DETECTION_DESTINATION_IP payload contains a hash of several pieces of information including the destination IP address and destination port value from the IP packet that is being sent to an IKE peer (for example, the packet containing the NAT_DETECTION_DESTINATION_IP payload).

Each IKEv2 peer sends one or more NAT_DETECTION_SOURCE_IP payloads and one NAT_DETECTION_DESTINATION_IP payload. The destination IP address and port of the outbound IKE packet are used to construct the hash that is contained within the first NAT_DETECTION_DESTINATION_IP payload. The source IP address and port of the outbound IKE packet are used to construct the hash that is contained within the NAT_DETECTION_SOURCE_IP payload. Normally, only one NAT_DETECTION_SOURCE_IP payload is sent; however, if the

sender of the packet has multiple IP addresses and it does not know which IP address is used to send the packet, it can send a NAT_DETECTION_SOURCE_IP payload for each IP address it owns.

z/OS provides limited support for IKEv2 NAT traversal, as well as additional z/OS-to-z/OS NAT traversal capabilities. A unique vendor ID is used to identify one z/OS IKE daemon to another. The z/OS IKE daemon sends a vendor ID payload incorporating the "z/OS CS-IKE NAT Traversal Level 1" vendor ID string shown in Table 108 on page 993 to identify itself as a z/OS IKE daemon to the peer. The z/OS IKE daemon includes this vendor ID payload in messages it sends that contain the NAT_DETECTION_SOURCE_IP and NAT_DETECTION_DESTINATION_IP payloads.

Impacts to IKEv1 phase 2 (Quick mode)

RFC 3947 defines two new encapsulation mode values: UDP-Encapsulated-Transport and UDP-Encapsulated-Tunnel. These new encapsulation modes are defined in RFC 3948. See z/OS Communications Server: IP Configuration Guide for a description of these new modes.

When one or more NAT devices are detected between two IKEv1 peers, messages 1 and 2 of a Quick mode exchange should not use offers containing tunnel or transport mode of encapsulation. Offers containing UDP-Encapsulated-Transport or UDP-Encapsulated-Tunnel mode of encapsulation should be used instead. Likewise, when no NAT devices are detected between two IKEv1 peers, messages 1 and 2 of a Quick mode exchange should not use offers containing UDP-Encapsulated-Transport or UDP-Encapsulated-Tunnel mode of encapsulation.

On z/OS, only the tunnel or transport mode of encapsulation can be specified on the IpDataOffer statement (see z/OS Communications Server: IP Configuration Reference). The decision to use UDP-Encapsulated-Transport or UDP-Encapsulated-Tunnel mode is made heuristically by the IKE daemon. When a NAT is detected between two IKEv1 peers, the z/OS IKE daemon converts IpDataOffer statements containing tunnel mode encapsulation to UDP-Encapsulated-Tunnel mode and IpDataOffers containing transport mode encapsulation to UDP-Encapsulated-Transport mode.

In order to facilitate incremental TCP and UDP checksum verification, RFC 3947 requires that IKEv1 peers exchange their view of each other's IP addresses when sending SA offers containing UDP-Encapsulated-Transport mode encapsulation. RFC 3947 defines a new payload for this purpose. This new payload is the NAT Original Address (NAT-OA) payload. The NAT-OA payload is an extension of RFC 2408 and 2409. It contains an IP address.

When the initiator of a Quick mode exchange sends a proposal utilizing UDP-Encapsulated-Transport mode, RFC 3947 requires the initiator to send two NAT-OA payload in message 1. The first NAT-OA payload contains the initiator's view of their IP address. The second NAT-OA payload contains the initiator's view of the responder's IP address.

When the responder of a Quick mode exchange accepts a proposal utilizing UDP-Encapsulated-Transport mode, RFC 3947 requires the responder to send two NAT-OA payloads in message 2. The first NAT-OA payload contains the responder's view of the initiator's address. The second NAT-OA payload contains the responder's view of his address.

In pre-RFC 3947 drafts, only one NAT-OA payload can be sent in messages 1 and 2 of a Quick mode exchange. Sending this NAT-OA payload was recommended when sending a proposal utilizing UDP-Encapsulated-Transport encapsulation, but not required. In message 1, it contained the initiator's view of his IP address. In message 2, it contained the responder's view of his IP address.

Impacts to IKEv2 phase 2

RFC 5996 defines only two encapsulation modes: tunnel mode and transport mode. On z/OS, tunnel or transport mode of encapsulation for IKEv2 can be specified on the HowToEncapIKEv2 parameter of the IpDynVpnAction statement (see z/OS Communications Server: IP Configuration Reference). The decision to use UDP-Encapsulated-Transport or UDP-Encapsulated-Tunnel mode is made heuristically by the IKE daemon. When a NAT is detected between two IKEv2 peers, a child SA that negotiates transport mode will use UDP-Encapsulated-Transport mode for the traffic protected by the child SA. When a NAT is detected between two IKEv2 peers, a child SA that does not negotiate transport mode will use UDP-Encapsulated-Tunnel mode for the traffic protected by the child SA. The TCP/IP stack must process received UDP-encapsulated ESP packets even when no NAT is detected.

In order to facilitate incremental TCP and UDP checksum verification when using transport mode encapsulation, the original source and destination IP address are obtained from the Traffic Selector payloads. When a NAT is detected and the initiator proposes transport mode, the Traffic Selector payloads may only contain one unique IP address, which is then used as the original IP address. A given Traffic Selector payload may contain multiple Traffic Selectors provided each Traffic Selector has the same IP address.

When a NAT has been detected and the initiator proposes transport mode, the request that includes the Traffic Selector payloads contains the initiator's view of the initiator's IP address in the TSi payload and includes the initiator's view of responder's IP address in the TSr payload. Likewise, if the responder accepts transport mode, the response that includes the Traffic Selector payloads includes the responder's view of the initiator's IP address in the TSi payload and includes the responder's view of responder's IP address in the TSr payload.

Utilizing port UDP 4500

To avoid any problems that could arise by IPSec-aware NAT devices, RFC 3947 and RFC 5996 both require the initiator to use UDP port 4500 to send and receive IKE traffic after the initiator detects the existence of a NAT device. RFC 3947 and RFC 5996 allow IKEv2 traffic to use port 4500 regardless of whether a NAT is detected, even when the initiator is sending the first phase 1 request. In Main mode, the initiator detects the existence of a NAT when processing message 4 and switches to source port UDP 4500 and destination port 4500 when the initiator is sending message 5. In Aggressive mode, the initiator detects the existence of a NAT when processing message 2 and switches to a source port of UDP 4500 and a destination port of UDP 4500 when sending message 3. In IKEv2, the initiator detects the existence of a NAT when processing the IKE_SA_INIT response. When the responder sends the initiator a message it must use the port values from the last message that was received from the initiator.

After the initiator switches to port 4500, which is known as port floating, all subsequent messages must use the floated ports. The initiator always

expects to send and receive messages on source port 4500 and destination port 4500. For the responder, if the remote peer is located behind a NAPT, the source port may have been changed to a value other than 4500. If so, the responder receives a message on a random source port Y and destination port 4500. After receiving this message, the responder sends subsequent messages using a source port of 4500 and destination port of Y.

These ports are also used to send UDP-encapsulated ESP traffic. In order to be able to distinguish UDP encapsulated ESP traffic from IKE traffic, a non-ESP marker is added to each IKE message sent using the UDP encapsulation ports. A non-ESP marker is 4 bytes of 0.

Figure 134 shows an IKE packet with and without the non-ESP marker.

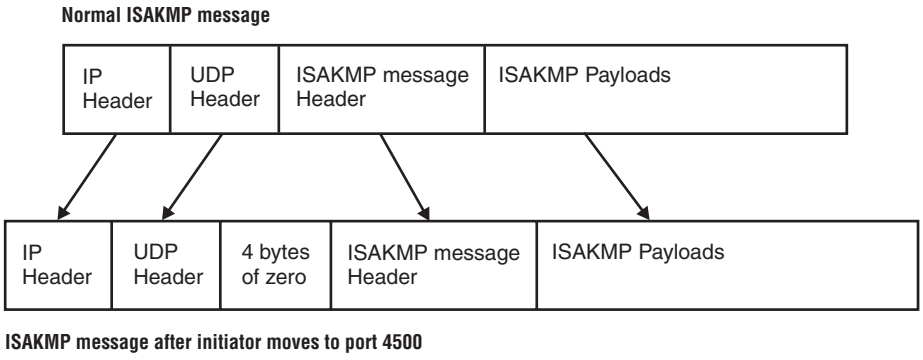

Figure 134. IKE packet with and without the non-ESP marker

Appendix C. 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.

For more information about IBM Health Checker for z/OS, see IBM Health Checker for z/OS: User's Guide.

z/OS Communications Server TCP/IP provides the following checks:

CSRES_AUTOQ_GLOBALTCPIPDATA
Checks whether the AUTOQUIESCE operand has been specified on the UNRESPONSIVETHRESHOLD resolver setup statement and that the GLOBALTCPIPDATA resolver setup statement has not been specified in the resolver setup file.

CSRES_AUTOQ_RESOLVEVIA
Checks whether the RESOLVEVIA statement has been specified with the value TCP in the global TCPIP.DATA file when the autonomic quiescing of unresponsive name servers function is active.

CSRES_AUTOQ_TIMEOUT
Checks whether the configured resolver timeout value in the global TCPIP.DATA file exceeds the optimal setting when the autonomic quiescing of unresponsive name servers function is active. By default, this check is performed once when the resolver is initialized and whenever a MODIFY REFRESH command is issued. This default value can be overridden on either a POLICY statement in the HZSPRM*xx* parmlib member or on a MODIFY command.

CSTCP_CINET_PORTRNG_RSV_*TCPIPstackname*
Checks whether the port range specified by INADDRANYPORT and INADDRANYCOUNT in the BPXPRMxx parmlib member is reserved for OMVS on this stack, when operating in a CINET environment. A port range is reserved on a TCP/IP stack using the PORTRANGE TCP/IP profile statement. By default, this check is performed once at stack initialization. This default can be overridden on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. The check name is suffixed by *TCPIPstackname*, which is the job name of each TCP/IP stack that is started, to define a separate check for each stack.

CSTCP_IPMAXRT4_*TCPIPstackname*
Checks whether the total number of IPv4 indirect routes in the TCP/IP stack routing table has exceeded the maximum threshold. When this threshold is exceeded, OMPROUTE and the TCP/IP stack can potentially experience high CPU consumption from routing changes. A large routing table is considered to be inefficient in network design and operation. By default, this check is performed at the following times:

- Whenever the total number of indirect routes exceeds the maximum threshold (default 2000)
- 30 minutes after stack initialization (provided that the maximum threshold has not been exceeded)

© Copyright IBM Corp. 2000, 2015

- Specified interval (default 168 hours for weekly)

The defaults for the maximum threshold and interval can be overridden on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. The check name is suffixed by *TCPIPstackname*, which is the job name of each TCP/IP stack that is started, to define a separate check for each stack.

CSTCP_IPMAXRT6_*TCPIPstackname*
Checks whether the total number of IPv6 indirect routes in the TCP/IP stack routing table has exceeded the maximum threshold. When this threshold is exceeded, OMPROUTE and the TCP/IP stack can potentially experience high CPU consumption from routing changes. A large routing table is considered to be inefficient in network design and operation. By default, this check is performed at the following times:

- Whenever the total number of indirect routes exceeds the maximum threshold (default 2000)
- 30 minutes after stack initialization (provided that the maximum threshold has not been exceeded)
- Specified interval (default 168 hours for weekly)

The defaults for the maximum threshold and interval can be overridden on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. The check name is suffixed by *TCPIPstackname*, which is the job name of each TCP/IP stack that is started, to define a separate check for each stack.

CSTCP_SYSTCPIP_CTRACE_*TCPIPstackname*
Checks whether TCP/IP Event Trace (SYSTCPIP) is active with options other than the default options (MINIMUM, INIT, OPCMDS, or OPMSGS). By default, this check will be performed once at stack initialization and then will be repeated once every 24 hours. This default can be overridden on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. The check name is suffixed by *TCPIPstackname*, which is the job name of each TCP stack that is started, to define a separate check for each stack.

CSTCP_SYSPLEXMON_RECOV_*TCPIPstackname*
Checks whether the IPCONFIG DYNAMICXCF or IPCONFIG6 DYNAMICXCF parameters have been specified and the GLOBALCONFIG SYSPLEXMONITOR RECOVERY parameter has been specified. This check produces an exception message if the IPCONFIG DYNAMICXCF or IPCONFIG6 DYNAMICXCF parameters were specified, but the GLOBALCONFIG SYSPLEXMONITOR NORECOVERY parameter is in effect. By default, this check is performed once at stack initialization. This default can be overridden on either a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. The check name is suffixed by *TCPIPstackname*, which is the job name of each TCP stack that is started, to define a separate check for each stack.

CSTCP_TCPMAXRCVBUFRSIZE_*TCPIPstackname*
Checks whether the configured TCP maximum receive buffer size is sufficient to provide optimal support to the z/OS Communications Server FTP Server. By default, this check is performed once at stack initialization and whenever a VARY TCPIP,,OBEYFILE command changes the TCPMAXRCVBUFRSIZE parameter. By default, it checks that TCPMAXRCVBUFRSIZE is at least 180K. These defaults can be overridden on either a POLICY statement in the HZSPRMxx parmlib member or on a

MODIFY command. The check name is suffixed by *TCPIPstackname*, which is the job name of each TCP stack that is started, to define a separate check for each stack.

ZOSMIGV2R2_Next_CS_SENDMAILDAEMN
Checks whether the sendmail daemon is in use on this system. By default, this check is inactive. This default can be overridden on a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. If an IBM Health Checker for z/OS exception message is generated, migration must be performed.

ZOSMIGV2R2_Next_CS_SENDMAILCLIEN
Checks whether the sendmail client is in use on this system. By default, this check is inactive. This default can be overridden on a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. If an IBM Health Checker for z/OS exception message is generated, migration must be performed.

ZOSMIGV2R2_Next_CS_SENDMAILMTA
Checks whether sendmail is used as a Mail Transfer Agent (MTA) listening on port 25 on this system. By default, this check is inactive. This default can be overridden on a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. If an IBM Health Checker for z/OS exception message is generated, migration must be performed.

ZOSMIGV2R2_Next_CS_SENDMAILMSA
Checks whether sendmail is used as a Mail Submission Agent (MSA) listening on port 587 on this system. By default, this check is inactive. This default can be overridden on a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. If an IBM Health Checker for z/OS exception message is generated, migration must be performed.

ZOSMIGV2R2_Next_CS_SMTPDDAEMON
Checks whether the SMTPD daemon is in use on this system. By default, this check is inactive. This default can be overridden on a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. If an IBM Health Checker for z/OS exception message is generated, migration must be performed.

ZOSMIGV2R2_Next_CS_SMTPDMTA
Checks whether the SMTPD daemon is in use as a Mail Transfer Agent (MTA) listening on port 25 on this system. By default, this check is inactive. This default can be overridden on a POLICY statement in the HZSPRMxx parmlib member or on a MODIFY command. If an IBM Health Checker for z/OS exception message is generated, migration must be performed.

Appendix D. Related protocol specifications

This appendix lists the related protocol specifications (RFCs) 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.

Draft RFCs that have been implemented in this and previous Communications Server releases are listed at the end of this topic.

Many features of TCP/IP Services are based on the following RFCs:

RFC Title and Author

RFC 652
Telnet output carriage-return disposition option D. Crocker

RFC 653
Telnet output horizontal tabstops option D. Crocker

RFC 654
Telnet output horizontal tab disposition option D. Crocker

RFC 655
Telnet output formfeed disposition option D. Crocker

RFC 657
Telnet output vertical tab disposition option D. Crocker

RFC 658
Telnet output linefeed disposition D. Crocker

RFC 698
Telnet extended ASCII option T. Mock

© Copyright IBM Corp. 2000, 2015

RFC 726
Remote Controlled Transmission and Echoing Telnet option J. Postel, D. Crocker

RFC 727
Telnet logout option M.R. Crispin

RFC 732
Telnet Data Entry Terminal option J.D. Day

RFC 733
Standard for the format of ARPA network text messages D. Crocker, J. Vittal, K.T. Pogran, D.A. Henderson

RFC 734
SUPDUP Protocol M.R. Crispin

RFC 735
Revised Telnet byte macro option D. Crocker, R.H. Gumpertz

RFC 736
Telnet SUPDUP option M.R. Crispin

RFC 749
Telnet SUPDUP—Output option B. Greenberg

RFC 765
File Transfer Protocol specification J. Postel

RFC 768
User Datagram Protocol J. Postel

RFC 779
Telnet send-location option E. Killian

RFC 783
TFTP Protocol (revision 2) K.R. Sollins

RFC 791
Internet Protocol J. Postel

RFC 792
Internet Control Message Protocol J. Postel

RFC 793
Transmission Control Protocol J. Postel

RFC 820
Assigned numbers J. Postel

RFC 821
Simple Mail Transfer Protocol J. Postel

RFC 822
Standard for the format of ARPA Internet text messages D. Crocker

RFC 823
DARPA Internet gateway R. Hinden, A. Sheltzer

RFC 826
Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware D. Plummer

RFC 854
Telnet Protocol Specification J. Postel, J. Reynolds

RFC 855
Telnet Option Specification J. Postel, J. Reynolds

RFC 856
Telnet Binary Transmission J. Postel, J. Reynolds

RFC 857
Telnet Echo Option J. Postel, J. Reynolds

RFC 858
Telnet Suppress Go Ahead Option J. Postel, J. Reynolds

RFC 859
Telnet Status Option J. Postel, J. Reynolds

RFC 860
Telnet Timing Mark Option J. Postel, J. Reynolds

RFC 861
Telnet Extended Options: List Option J. Postel, J. Reynolds

RFC 862
Echo Protocol J. Postel

RFC 863
Discard Protocol J. Postel

RFC 864
Character Generator Protocol J. Postel

RFC 865
Quote of the Day Protocol J. Postel

RFC 868
Time Protocol J. Postel, K. Harrenstien

RFC 877
Standard for the transmission of IP datagrams over public data networks J.T. Korb

RFC 883
Domain names: Implementation specification P.V. Mockapetris

RFC 884
Telnet terminal type option M. Solomon, E. Wimmers

RFC 885
Telnet end of record option J. Postel

RFC 894
Standard for the transmission of IP datagrams over Ethernet networks C. Hornig

RFC 896
Congestion control in IP/TCP internetworks J. Nagle

RFC 903
Reverse Address Resolution Protocol R. Finlayson, T. Mann, J. Mogul, M. Theimer

RFC 904
Exterior Gateway Protocol formal specification D. Mills

RFC 919
Broadcasting Internet Datagrams J. Mogul

RFC 922
Broadcasting Internet datagrams in the presence of subnets J. Mogul

RFC 927
TACACS user identification Telnet option B.A. Anderson

RFC 933
Output marking Telnet option S. Silverman

RFC 946
Telnet terminal location number option R. Nedved

RFC 950
Internet Standard Subnetting Procedure J. Mogul, J. Postel

RFC 952
DoD Internet host table specification K. Harrenstien, M. Stahl, E. Feinler

RFC 959
File Transfer Protocol J. Postel, J.K. Reynolds

RFC 961
Official ARPA-Internet protocols J.K. Reynolds, J. Postel

RFC 974
Mail routing and the domain system C. Partridge

RFC 1001
Protocol standard for a NetBIOS service on a TCP/UDP transport: Concepts and methods NetBios Working Group in the Defense Advanced Research Projects Agency, Internet Activities Board, End-to-End Services Task Force

RFC 1002
Protocol Standard for a NetBIOS service on a TCP/UDP transport: Detailed specifications NetBios Working Group in the Defense Advanced Research Projects Agency, Internet Activities Board, End-to-End Services Task Force

RFC 1006
ISO transport services on top of the TCP: Version 3 M.T. Rose, D.E. Cass

RFC 1009
Requirements for Internet gateways R. Braden, J. Postel

RFC 1011
Official Internet protocols J. Reynolds, J. Postel

RFC 1013
X Window System Protocol, version 11: Alpha update April 1987 R. Scheifler

RFC 1014
XDR: External Data Representation standard Sun Microsystems

RFC 1027
Using ARP to implement transparent subnet gateways S. Carl-Mitchell, J. Quarterman

RFC 1032
Domain administrators guide M. Stahl

RFC 1033
Domain administrators operations guide M. Lottor

RFC 1034
Domain names—concepts and facilities P.V. Mockapetris

RFC 1035
Domain names—implementation and specification P.V. Mockapetris

RFC 1038
Draft revised IP security option M. St. Johns

RFC 1041
Telnet 3270 regime option Y. Rekhter

RFC 1042
Standard for the transmission of IP datagrams over IEEE 802 networks J. Postel, J. Reynolds

RFC 1043
Telnet Data Entry Terminal option: DODIIS implementation A. Yasuda, T. Thompson

RFC 1044
Internet Protocol on Network System's HYPERchannel: Protocol specification K. Hardwick, J. Lekashman

RFC 1053
Telnet X.3 PAD option S. Levy, T. Jacobson

RFC 1055
Nonstandard for transmission of IP datagrams over serial lines: SLIP J. Romkey

RFC 1057
RPC: Remote Procedure Call Protocol Specification: Version 2 Sun Microsystems

RFC 1058
Routing Information Protocol C. Hedrick

RFC 1060
Assigned numbers J. Reynolds, J. Postel

RFC 1067
Simple Network Management Protocol J.D. Case, M. Fedor, M.L. Schoffstall, J. Davin

RFC 1071
Computing the Internet checksum R.T. Braden, D.A. Borman, C. Partridge

RFC 1072
TCP extensions for long-delay paths V. Jacobson, R.T. Braden

RFC 1073
Telnet window size option D. Waitzman

RFC 1079
Telnet terminal speed option C. Hedrick

RFC 1085
ISO presentation services on top of TCP/IP based internets M.T. Rose

RFC 1091
Telnet terminal-type option J. VanBokkelen

RFC 1094
NFS: Network File System Protocol specification Sun Microsystems

RFC 1096
Telnet X display location option G. Marcy

RFC 1101
DNS encoding of network names and other types P. Mockapetris

RFC 1112
Host extensions for IP multicasting S.E. Deering

RFC 1113
Privacy enhancement for Internet electronic mail: Part I — message encipherment and authentication procedures J. Linn

RFC 1118
Hitchhikers Guide to the Internet E. Krol

RFC 1122
Requirements for Internet Hosts—Communication Layers R. Braden, Ed.

RFC 1123
Requirements for Internet Hosts—Application and Support R. Braden, Ed.

RFC 1146
TCP alternate checksum options J. Zweig, C. Partridge

RFC 1155
Structure and identification of management information for TCP/IP-based internets M. Rose, K. McCloghrie

RFC 1156
Management Information Base for network management of TCP/IP-based internets K. McCloghrie, M. Rose

RFC 1157
Simple Network Management Protocol (SNMP) J. Case, M. Fedor, M. Schoffstall, J. Davin

RFC 1158
Management Information Base for network management of TCP/IP-based internets: MIB-II M. Rose

RFC 1166
Internet numbers S. Kirkpatrick, M.K. Stahl, M. Recker

RFC 1179
Line printer daemon protocol L. McLaughlin

RFC 1180
TCP/IP tutorial T. Socolofsky, C. Kale

RFC 1183
New DNS RR Definitions C.F. Everhart, L.A. Mamakos, R. Ullmann, P.V. Mockapetris

RFC 1184
Telnet Linemode Option D. Borman

RFC 1186
MD4 Message Digest Algorithm R.L. Rivest

RFC 1187
Bulk Table Retrieval with the SNMP M. Rose, K. McCloghrie, J. Davin

RFC 1188
Proposed Standard for the Transmission of IP Datagrams over FDDI Networks D. Katz

RFC 1190
Experimental Internet Stream Protocol: Version 2 (ST-II) C. Topolcic

RFC 1191
Path MTU discovery J. Mogul, S. Deering

RFC 1198
FYI on the X window system R. Scheifler

RFC 1207
FYI on Questions and Answers: Answers to commonly asked "experienced Internet user" questions G. Malkin, A. Marine, J. Reynolds

RFC 1208
Glossary of networking terms O. Jacobsen, D. Lynch

RFC 1213
Management Information Base for Network Management of TCP/IP-based internets: MIB-II K. McCloghrie, M.T. Rose

RFC 1215
Convention for defining traps for use with the SNMP M. Rose

RFC 1227
SNMP MUX protocol and MIB M.T. Rose

RFC 1228
SNMP-DPI: Simple Network Management Protocol Distributed Program Interface G. Carpenter, B. Wijnen

RFC 1229
Extensions to the generic-interface MIB K. McCloghrie

RFC 1230
IEEE 802.4 Token Bus MIB K. McCloghrie, R. Fox

RFC 1231
IEEE 802.5 Token Ring MIB K. McCloghrie, R. Fox, E. Decker

RFC 1236
IP to X.121 address mapping for DDN L. Morales, P. Hasse

RFC 1256
ICMP Router Discovery Messages S. Deering, Ed.

RFC 1267
Border Gateway Protocol 3 (BGP-3) K. Lougheed, Y. Rekhter

RFC 1268
Application of the Border Gateway Protocol in the Internet Y. Rekhter, P. Gross

RFC 1269
Definitions of Managed Objects for the Border Gateway Protocol: Version 3 S. Willis, J. Burruss

RFC 1270
SNMP Communications Services F. Kastenholz, ed.

RFC 1285
FDDI Management Information Base J. Case

RFC 1315
Management Information Base for Frame Relay DTEs C. Brown, F. Baker, C. Carvalho

RFC 1321
The MD5 Message-Digest Algorithm R. Rivest

RFC 1323
TCP Extensions for High Performance V. Jacobson, R. Braden, D. Borman

RFC 1325
FYI on Questions and Answers: Answers to Commonly Asked "New Internet User" Questions G. Malkin, A. Marine

RFC 1327
Mapping between X.400 (1988)/ISO 10021 and RFC 822 S. Hardcastle-Kille

RFC 1340
Assigned Numbers J. Reynolds, J. Postel

RFC 1344
Implications of MIME for Internet Mail Gateways N. Bornstein

RFC 1349
Type of Service in the Internet Protocol Suite P. Almquist

RFC 1350
The TFTP Protocol (Revision 2) K.R. Sollins

RFC 1351
SNMP Administrative Model J. Davin, J. Galvin, K. McCloghrie

RFC 1352
SNMP Security Protocols J. Galvin, K. McCloghrie, J. Davin

RFC 1353
Definitions of Managed Objects for Administration of SNMP Parties K. McCloghrie, J. Davin, J. Galvin

RFC 1354
IP Forwarding Table MIB F. Baker

RFC 1356
Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode A. Malis, D. Robinson, R. Ullmann

RFC 1358
Charter of the Internet Architecture Board (IAB) L. Chapin

RFC 1363
A Proposed Flow Specification C. Partridge

RFC 1368
Definition of Managed Objects for IEEE 802.3 Repeater Devices D. McMaster, K. McCloghrie

RFC 1372
Telnet Remote Flow Control Option C. L. Hedrick, D. Borman

RFC 1374
IP and ARP on HIPPI J. Renwick, A. Nicholson

RFC 1381
SNMP MIB Extension for X.25 LAPB D. Throop, F. Baker

RFC 1382
SNMP MIB Extension for the X.25 Packet Layer D. Throop

RFC 1387
RIP Version 2 Protocol Analysis G. Malkin

RFC 1388
RIP Version 2 Carrying Additional Information G. Malkin

RFC 1389
RIP Version 2 MIB Extensions G. Malkin, F. Baker

RFC 1390
Transmission of IP and ARP over FDDI Networks D. Katz

RFC 1393
Traceroute Using an IP Option G. Malkin

RFC 1398
Definitions of Managed Objects for the Ethernet-Like Interface Types F. Kastenholz

RFC 1408
Telnet Environment Option D. Borman, Ed.

RFC 1413
Identification Protocol M. St. Johns

RFC 1416
Telnet Authentication Option D. Borman, ed.

RFC 1420
SNMP over IPX S. Bostock

RFC 1428
Transition of Internet Mail from Just-Send-8 to 8bit-SMTP/MIME G. Vaudreuil

RFC 1442
Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1443
Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1445
Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2) J. Galvin, K. McCloghrie

RFC 1447
Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2) 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 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 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 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
SMTP Service Extension for 8bit-MIMEtransport J. Klensin, N. Freed, M. Rose, E. Stefferud, D. Crocker

RFC 1664
Using the Internet DNS to Distribute RFC1327 Mail Address Mapping Tables C. Allochio, A. Bonito, B. Cole, S. Giordano, R. Hagens

RFC 1693
An Extension to TCP: Partial Order Service T. Connolly, P. Amer, P. Conrad

RFC 1695
Definitions of Managed Objects for ATM Management Version 8.0 using SMIv2 M. Ahmed, K. Tesink

RFC 1701
Generic Routing Encapsulation (GRE) S. Hanks, T. Li, D. Farinacci, P. Traina

RFC 1702
Generic Routing Encapsulation over IPv4 networks S. Hanks, T. Li, D. Farinacci, P. Traina

RFC 1706
DNS NSAP Resource Records B. Manning, R. Colella

RFC 1712
DNS Encoding of Geographical Location C. Farrell, M. Schulze, S. Pleitner D. Baldoni

RFC 1713
Tools for DNS debugging A. Romao

RFC 1723
RIP Version 2—Carrying Additional Information G. Malkin

RFC 1752
The Recommendation for the IP Next Generation Protocol S. Bradner, A. Mankin

RFC 1766
Tags for the Identification of Languages H. Alvestrand

RFC 1771
A Border Gateway Protocol 4 (BGP-4) Y. Rekhter, T. Li

RFC 1794
DNS Support for Load Balancing T. Brisco

RFC 1819
Internet Stream Protocol Version 2 (ST2) Protocol Specification—Version ST2+ L. Delgrossi, L. Berger Eds.

RFC 1826
IP Authentication Header R. Atkinson

RFC 1828
IP Authentication using Keyed MD5 P. Metzger, W. Simpson

RFC 1829
The ESP DES-CBC Transform P. Karn, P. Metzger, W. Simpson

RFC 1830
SMTP Service Extensions for Transmission of Large and Binary MIME Messages G. Vaudreuil

RFC 1831
RPC: Remote Procedure Call Protocol Specification Version 2 R. Srinivasan

RFC 1832
XDR: External Data Representation Standard R. Srinivasan

RFC 1833
Binding Protocols for ONC RPC Version 2 R. Srinivasan

RFC 1850
OSPF Version 2 Management Information Base F. Baker, R. Coltun

RFC 1854
SMTP Service Extension for Command Pipelining N. Freed

RFC 1869
SMTP Service Extensions J. Klensin, N. Freed, M. Rose, E. Stefferud, D. Crocker

RFC 1870
SMTP Service Extension for Message Size Declaration J. Klensin, N. Freed, K. Moore

RFC 1876
A Means for Expressing Location Information in the Domain Name System C. Davis, P. Vixie, T. Goodwin, I. Dickinson

RFC 1883
Internet Protocol, Version 6 (IPv6) Specification S. Deering, R. Hinden

RFC 1884
IP Version 6 Addressing Architecture R. Hinden, S. Deering, Eds.

RFC 1886
DNS Extensions to support IP version 6 S. Thomson, C. Huitema

RFC 1888
OSI NSAPs and IPv6 J. Bound, B. Carpenter, D. Harrington, J. Houldsworth, A. Lloyd

RFC 1891
SMTP Service Extension for Delivery Status Notifications K. Moore

RFC 1892
The Multipart/Report Content Type for the Reporting of Mail System Administrative Messages G. Vaudreuil

RFC 1894
*An Extensible Message Format for Delivery Status Notifications*K. Moore, G. Vaudreuil

RFC 1901
Introduction to Community-based SNMPv2 J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1902
Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1903
Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1904
Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1905
Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1906
Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1907
Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2) J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1908
Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 1912
Common DNS Operational and Configuration Errors D. Barr

RFC 1918
Address Allocation for Private Internets Y. Rekhter, B. Moskowitz, D. Karrenberg, G.J. de Groot, E. Lear

RFC 1928
SOCKS Protocol Version 5 M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, L. Jones

RFC 1930
Guidelines for creation, selection, and registration of an Autonomous System (AS) J. Hawkinson, T. Bates

RFC 1939
Post Office Protocol-Version 3 J. Myers, M. Rose

RFC 1981
Path MTU Discovery for IP version 6 J. McCann, S. Deering, J. Mogul

RFC 1982
Serial Number Arithmetic R. Elz, R. Bush

RFC 1985
SMTP Service Extension for Remote Message Queue Starting J. De Winter

RFC 1995
Incremental Zone Transfer in DNS M. Ohta

RFC 1996
A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY) P. Vixie

RFC 2010
Operational Criteria for Root Name Servers B. Manning, P. Vixie

RFC 2011
SNMPv2 Management Information Base for the Internet Protocol using SMIv2 K. McCloghrie, Ed.

RFC 2012
SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2 K. McCloghrie, Ed.

RFC 2013
SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2 K. McCloghrie, Ed.

RFC 2018
TCP Selective Acknowledgement Options M. Mathis, J. Mahdavi, S. Floyd, A. Romanow

RFC 2026
The Internet Standards Process — Revision 3 S. Bradner

RFC 2030
Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI D. Mills

RFC 2033
Local Mail Transfer Protocol J. Myers

RFC 2034
*SMTP Service Extension for Returning Enhanced Error Codes*N. Freed

RFC 2040
*The RC5, RC5–CBC, RC-5–CBC-Pad, and RC5–CTS Algorithms*R. Baldwin, R. Rivest

RFC 2045
Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies N. Freed, N. Borenstein

RFC 2052
A DNS RR for specifying the location of services (DNS SRV) A. Gulbrandsen, P. Vixie

RFC 2065
Domain Name System Security Extensions D. Eastlake 3rd, C. Kaufman

RFC 2066
TELNET CHARSET Option R. Gellens

RFC 2080
RIPng for IPv6 G. Malkin, R. Minnear

RFC 2096
IP Forwarding Table MIB F. Baker

RFC 2104
HMAC: Keyed-Hashing for Message Authentication H. Krawczyk, M. Bellare, R. Canetti

RFC 2119
Keywords for use in RFCs to Indicate Requirement Levels S. Bradner

RFC 2133
Basic Socket Interface Extensions for IPv6 R. Gilligan, S. Thomson, J. Bound, W. Stevens

RFC 2136
Dynamic Updates in the Domain Name System (DNS UPDATE) P. Vixie, Ed., S. Thomson, Y. Rekhter, J. Bound

RFC 2137
Secure Domain Name System Dynamic Update D. Eastlake 3rd

RFC 2163
Using the Internet DNS to Distribute MIXER Conformant Global Address Mapping (MCGAM) C. Allocchio

RFC 2168
Resolution of Uniform Resource Identifiers using the Domain Name System R. Daniel, M. Mealling

RFC 2178
OSPF Version 2 J. Moy

RFC 2181
Clarifications to the DNS Specification R. Elz, R. Bush

RFC 2205
Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, S. Jamin

RFC 2210
The Use of RSVP with IETF Integrated Services J. Wroclawski

RFC 2211
Specification of the Controlled-Load Network Element Service J. Wroclawski

RFC 2212
Specification of Guaranteed Quality of Service S. Shenker, C. Partridge, R. Guerin

RFC 2215
General Characterization Parameters for Integrated Service Network Elements S. Shenker, J. Wroclawski

RFC 2217
Telnet Com Port Control Option G. Clarke

RFC 2219
Use of DNS Aliases for Network Services M. Hamilton, R. Wright

RFC 2228
FTP Security Extensions M. Horowitz, S. Lunt

RFC 2230
Key Exchange Delegation Record for the DNS R. Atkinson

RFC 2233
The Interfaces Group MIB using SMIv2 K. McCloghrie, F. Kastenholz

RFC 2240
A Legal Basis for Domain Name Allocation O. Vaughn

RFC 2246
The TLS Protocol Version 1.0 T. Dierks, C. Allen

RFC 2251
Lightweight Directory Access Protocol (v3) M. Wahl, T. Howes, S. Kille

RFC 2253
Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names M. Wahl, S. Kille, T. Howes

RFC 2254
The String Representation of LDAP Search Filters T. Howes

RFC 2261
An Architecture for Describing SNMP Management Frameworks D. Harrington, R. Presuhn, B. Wijnen

RFC 2262
Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) J. Case, D. Harrington, R. Presuhn, B. Wijnen

RFC 2271
An Architecture for Describing SNMP Management Frameworks D. Harrington, R. Presuhn, B. Wijnen

RFC 2273
SNMPv3 Applications D. Levi, P. Meyer, B. Stewartz

RFC 2274
User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) U. Blumenthal, B. Wijnen

RFC 2275
View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) B. Wijnen, R. Presuhn, K. McCloghrie

RFC 2279
UTF-8, a transformation format of ISO 10646 F. Yergeau

RFC 2292
Advanced Sockets API for IPv6 W. Stevens, M. Thomas

RFC 2308
Negative Caching of DNS Queries (DNS NCACHE) M. Andrews

RFC 2317
Classless IN-ADDR.ARPA delegation H. Eidnes, G. de Groot, P. Vixie

RFC 2320
Definitions of Managed Objects for Classical IP and ARP Over ATM Using SMIv2 (IPOA-MIB) M. Greene, J. Luciani, K. White, T. Kuo

RFC 2328
OSPF Version 2 J. Moy

RFC 2345
Domain Names and Company Name Retrieval J. Klensin, T. Wolf, G. Oglesby

RFC 2352
A Convention for Using Legal Names as Domain Names O. Vaughn

RFC 2355
TN3270 Enhancements B. Kelly

RFC 2358
Definitions of Managed Objects for the Ethernet-like Interface Types J. Flick, J. Johnson

RFC 2373
IP Version 6 Addressing Architecture R. Hinden, S. Deering

RFC 2374
An IPv6 Aggregatable Global Unicast Address Format R. Hinden, M. O'Dell, S. Deering

RFC 2375
IPv6 Multicast Address Assignments R. Hinden, S. Deering

RFC 2385
Protection of BGP Sessions via the TCP MD5 Signature Option A. Hefferman

RFC 2389
Feature negotiation mechanism for the File Transfer Protocol P. Hethmon, R. Elz

RFC 2401
Security Architecture for Internet Protocol S. Kent, R. Atkinson

RFC 2402
IP Authentication Header S. Kent, R. Atkinson

RFC 2403
The Use of HMAC-MD5–96 within ESP and AH C. Madson, R. Glenn

RFC 2404
The Use of HMAC-SHA–1–96 within ESP and AH C. Madson, R. Glenn

RFC 2405
The ESP DES-CBC Cipher Algorithm With Explicit IV C. Madson, N. Doraswamy

RFC 2406
IP Encapsulating Security Payload (ESP) S. Kent, R. Atkinson

RFC 2407
*The Internet IP Security Domain of Interpretation for ISAKMP*D. Piper

RFC 2408
Internet Security Association and Key Management Protocol (ISAKMP) D. Maughan, M. Schertler, M. Schneider, J. Turner

RFC 2409
The Internet Key Exchange (IKE) D. Harkins, D. Carrel

RFC 2410
The NULL Encryption Algorithm and Its Use With IPsec R. Glenn, S. Kent,

RFC 2428
FTP Extensions for IPv6 and NATs M. Allman, S. Ostermann, C. Metz

RFC 2445
Internet Calendaring and Scheduling Core Object Specification (iCalendar) F. Dawson, D. Stenerson

RFC 2459
Internet X.509 Public Key Infrastructure Certificate and CRL Profile R. Housley, W. Ford, W. Polk, D. Solo

RFC 2460
Internet Protocol, Version 6 (IPv6) Specification S. Deering, R. Hinden

RFC 2461
Neighbor Discovery for IP Version 6 (IPv6) T. Narten, E. Nordmark, W. Simpson

RFC 2462
IPv6 Stateless Address Autoconfiguration S. Thomson, T. Narten

RFC 2463
Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification A. Conta, S. Deering

RFC 2464
Transmission of IPv6 Packets over Ethernet Networks M. Crawford

RFC 2466
Management Information Base for IP Version 6: ICMPv6 Group D. Haskin, S. Onishi

RFC 2476
Message Submission R. Gellens, J. Klensin

RFC 2487
SMTP Service Extension for Secure SMTP over TLS P. Hoffman

RFC 2505
Anti-Spam Recommendations for SMTP MTAs G. Lindberg

RFC 2523
Photuris: Extended Schemes and Attributes P. Karn, W. Simpson

RFC 2535
Domain Name System Security Extensions D. Eastlake 3rd

RFC 2538
Storing Certificates in the Domain Name System (DNS) D. Eastlake 3rd, O. Gudmundsson

RFC 2539
Storage of Diffie-Hellman Keys in the Domain Name System (DNS) D. Eastlake 3rd

RFC 2540
Detached Domain Name System (DNS) Information D. Eastlake 3rd

RFC 2554
SMTP Service Extension for Authentication J. Myers

RFC 2570
Introduction to Version 3 of the Internet-standard Network Management Framework J. Case, R. Mundy, D. Partain, B. Stewart

RFC 2571
An Architecture for Describing SNMP Management Frameworks B. Wijnen, D. Harrington, R. Presuhn

RFC 2572
Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) J. Case, D. Harrington, R. Presuhn, B. Wijnen

RFC 2573
SNMP Applications D. Levi, P. Meyer, B. Stewart

RFC 2574
User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) U. Blumenthal, B. Wijnen

RFC 2575
View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) B. Wijnen, R. Presuhn, K. McCloghrie

RFC 2576
Co-Existence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework R. Frye, D. Levi, S. Routhier, B. Wijnen

RFC 2578
Structure of Management Information Version 2 (SMIv2) K. McCloghrie, D. Perkins, J. Schoenwaelder

RFC 2579
Textual Conventions for SMIv2 K. McCloghrie, D. Perkins, J. Schoenwaelder

RFC 2580
Conformance Statements for SMIv2 K. McCloghrie, D. Perkins, J. Schoenwaelder

RFC 2581
TCP Congestion Control M. Allman, V. Paxson, W. Stevens

RFC 2583
Guidelines for Next Hop Client (NHC) Developers R. Carlson, L. Winkler

RFC 2591
Definitions of Managed Objects for Scheduling Management Operations D. Levi, J. Schoenwaelder

RFC 2625
IP and ARP over Fibre Channel M. Rajagopal, R. Bhagwat, W. Rickard

RFC 2635
Don't SPEW A Set of Guidelines for Mass Unsolicited Mailings and Postings (spam)* S. Hambridge, A. Lunde

RFC 2637
Point-to-Point Tunneling Protocol K. Hamzeh, G. Pall, W. Verthein, J. Taarud, W. Little, G. Zorn

RFC 2640
Internationalization of the File Transfer Protocol B. Curtin

RFC 2665
Definitions of Managed Objects for the Ethernet-like Interface Types J. Flick, J. Johnson

RFC 2671
Extension Mechanisms for DNS (EDNS0) P. Vixie

RFC 2672
Non-Terminal DNS Name Redirection M. Crawford

RFC 2675
IPv6 Jumbograms D. Borman, S. Deering, R. Hinden

RFC 2710
Multicast Listener Discovery (MLD) for IPv6 S. Deering, W. Fenner, B. Haberman

RFC 2711
IPv6 Router Alert Option C. Partridge, A. Jackson

RFC 2740
OSPF for IPv6 R. Coltun, D. Ferguson, J. Moy

RFC 2753
A Framework for Policy-based Admission Control R. Yavatkar, D. Pendarakis, R. Guerin

RFC 2782
A DNS RR for specifying the location of services (DNS SRV) A. Gubrandsen, P. Vixix, L. Esibov

RFC 2821
Simple Mail Transfer Protocol J. Klensin, Ed.

RFC 2822
Internet Message Format P. Resnick, Ed.

RFC 2840
TELNET KERMIT OPTION J. Altman, F. da Cruz

RFC 2845
Secret Key Transaction Authentication for DNS (TSIG) P. Vixie, O. Gudmundsson, D. Eastlake 3rd, B. Wellington

RFC 2851
Textual Conventions for Internet Network Addresses M. Daniele, B. Haberman, S. Routhier, J. Schoenwaelder

RFC 2852
Deliver By SMTP Service Extension D. Newman

RFC 2874
DNS Extensions to Support IPv6 Address Aggregation and Renumbering M. Crawford, C. Huitema

RFC 2915
The Naming Authority Pointer (NAPTR) DNS Resource Record M. Mealling, R. Daniel

RFC 2920
SMTP Service Extension for Command Pipelining N. Freed

RFC 2930
Secret Key Establishment for DNS (TKEY RR) D. Eastlake, 3rd

RFC 2941
Telnet Authentication Option T. Ts'o, ed., J. Altman

RFC 2942
Telnet Authentication: Kerberos Version 5 T. Ts'o

RFC 2946
Telnet Data Encryption Option T. Ts'o

RFC 2952
Telnet Encryption: DES 64 bit Cipher Feedback T. Ts'o

RFC 2953
Telnet Encryption: DES 64 bit Output Feedback T. Ts'o

RFC 2992
Analysis of an Equal-Cost Multi-Path Algorithm C. Hopps

RFC 3019
IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol B. Haberman, R. Worzella

RFC 3060
Policy Core Information Model—Version 1 Specification B. Moore, E. Ellesson, J. Strassner, A. Westerinen

RFC 3152
Delegation of IP6.ARPA R. Bush

RFC 3164
The BSD Syslog Protocol C. Lonvick

RFC 3207
SMTP Service Extension for Secure SMTP over Transport Layer Security P. Hoffman

RFC 3226
DNSSEC and IPv6 A6 aware server/resolver message size requirements O. Gudmundsson

RFC 3291
Textual Conventions for Internet Network Addresses M. Daniele, B. Haberman, S. Routhier, J. Schoenwaelder

RFC 3363
Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System R. Bush, A. Durand, B. Fink, O. Gudmundsson, T. Hain

RFC 3376
Internet Group Management Protocol, Version 3 B. Cain, S. Deering, I. Kouvelas, B. Fenner, A. Thyagarajan

RFC 3390
Increasing TCP's Initial Window M. Allman, S. Floyd, C. Partridge

RFC 3410
Introduction and Applicability Statements for Internet-Standard Management Framework J. Case, R. Mundy, D. Partain, B. Stewart

RFC 3411
An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks D. Harrington, R. Presuhn, B. Wijnen

RFC 3412
Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) J. Case, D. Harrington, R. Presuhn, B. Wijnen

RFC 3413
Simple Network Management Protocol (SNMP) Applications D. Levi, P. Meyer, B. Stewart

RFC 3414
User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) U. Blumenthal, B. Wijnen

RFC 3415
View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) B. Wijnen, R. Presuhn, K. McCloghrie

RFC 3416
Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) R. Presuhn, J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 3417
Transport Mappings for the Simple Network Management Protocol (SNMP) R. Presuhn, J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 3418
Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) R. Presuhn, J. Case, K. McCloghrie, M. Rose, S. Waldbusser

RFC 3419
Textual Conventions for Transport Addresses M. Daniele, J. Schoenwaelder

RFC 3484
Default Address Selection for Internet Protocol version 6 (IPv6) R. Draves

RFC 3493
Basic Socket Interface Extensions for IPv6 R. Gilligan, S. Thomson, J. Bound, J. McCann, W. Stevens

RFC 3513
Internet Protocol Version 6 (IPv6) Addressing Architecture R. Hinden, S. Deering

RFC 3526
More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE) T. Kivinen, M. Kojo

RFC 3542
Advanced Sockets Application Programming Interface (API) for IPv6 W. Richard Stevens, M. Thomas, E. Nordmark, T. Jinmei

RFC 3566
The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec S. Frankel, H. Herbert

RFC 3569
An Overview of Source-Specific Multicast (SSM) S. Bhattacharyya, Ed.

RFC 3584
Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework R. Frye, D. Levi, S. Routhier, B. Wijnen

RFC 3602
The AES-CBC Cipher Algorithm and Its Use with IPsec S. Frankel, R. Glenn, S. Kelly

RFC 3629
UTF-8, a transformation format of ISO 10646 R. Kermode, C. Vicisano

RFC 3658
Delegation Signer (DS) Resource Record (RR) O. Gudmundsson

RFC 3678
Socket Interface Extensions for Multicast Source Filters D. Thaler, B. Fenner, B. Quinn

RFC 3715
IPsec-Network Address Translation (NAT) Compatibility Requirements B. Aboba, W. Dixon

RFC 3810
Multicast Listener Discovery Version 2 (MLDv2) for IPv6 R. Vida, Ed., L. Costa, Ed.

RFC 3826
The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model U. Blumenthal, F. Maino, K McCloghrie.

RFC 3947
Negotiation of NAT-Traversal in the IKE T. Kivinen, B. Swander, A. Huttunen, V. Volpe

RFC 3948
UDP Encapsulation of IPsec ESP Packets A. Huttunen, B. Swander, V. Volpe, L. DiBurro, M. Stenberg

RFC 4001
Textual Conventions for Internet Network Addresses M. Daniele, B. Haberman, S. Routhier, J. Schoenwaelder

RFC 4007
IPv6 Scoped Address Architecture S. Deering, B. Haberman, T. Jinmei, E. Nordmark, B. Zill

RFC 4022
Management Information Base for the Transmission Control Protocol (TCP) R. Raghunarayan

RFC 4106
The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP) J. Viega, D. McGrew

RFC 4109
Algorithms for Internet Key Exchange version 1 (IKEv1) P. Hoffman

RFC 4113
Management Information Base for the User Datagram Protocol (UDP) B. Fenner, J. Flick

RFC 4191
Default Router Preferences and More-Specific Routes R. Draves, D. Thaler

RFC 4217
Securing FTP with TLS P. Ford-Hutchinson

RFC 4292
IP Forwarding Table MIB B. Haberman

RFC 4293
Management Information Base for the Internet Protocol (IP) S. Routhier

RFC 4301
Security Architecture for the Internet Protocol S. Kent, K. Seo

RFC 4302
IP Authentication Header S. Kent

RFC 4303
IP Encapsulating Security Payload (ESP) S. Kent

RFC 4304
Extended Sequence Number (ESN) Addendum to IPsec Domain of Interpretation (DOI) for Internet Security Association and Key Management Protocol (ISAKMP) S. Kent

RFC 4307
Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2) J. Schiller

RFC 4308
Cryptographic Suites for IPsec P. Hoffman

RFC 4434
The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol P. Hoffman

RFC 4443
Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification A. Conta, S. Deering

RFC 4552
Authentication/Confidentiality for OSPFv3 M. Gupta, N. Melam

RFC 4678
Server/Application State Protocol v1 A. Bivens

RFC 4753
ECP Groups for IKE and IKEv2 D. Fu, J. Solinas

RFC 4754
IKE and IKEv2 Authentication Using the Elliptic Curve Digital Signature Algorithm (ECDSA) D. Fu, J. Solinas

RFC 4809
Requirements for an IPsec Certificate Management Profile C. Bonatti, Ed., S. Turner, Ed., G. Lebovitz, Ed.

RFC 4835
Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH) V. Manral

RFC 4862
IPv6 Stateless Address Autoconfiguration S. Thomson, T. Narten, T. Jinmei

RFC 4868
Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec S. Kelly, S. Frankel

RFC 4869
Suite B Cryptographic Suites for IPsec L. Law, J. Solinas

RFC 4941
Privacy Extensions for Stateless Address Autoconfiguration in IPv6 T. Narten, R. Draves, S. Krishnan

RFC 4945
The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2, and PKIX B. Korver

RFC 5014
IPv6 Socket API for Source Address Selection E. Nordmark, S. Chakrabarti, J. Laganier

RFC 5095
Deprecation of Type 0 Routing Headers in IPv6 J. Abley, P. Savola, G. Neville-Neil

RFC 5175
IPv6 Router Advertisement Flags Option B. Haberman, Ed., R. Hinden

RFC 5282
Using Authenticated Encryption Algorithms with the Encrypted Payload of the Internet Key Exchange version 2 (IKEv2) Protocol D. Black, D. McGrew

RFC 5996
Internet Key Exchange Protocol Version 2 (IKEv2) C. Kaufman, P. Hoffman, Y. Nir, P. Eronen

Internet drafts

Internet drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Other groups can also distribute working documents as Internet drafts. You can see Internet drafts at http://www.ietf.org/ID.html.

Appendix E. Accessibility

Publications for this product are offered in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties when using PDF files, you can view the information through the z/OS Internet Library website or IBM Knowledge Center. If you continue to experience problems, send an email to mhvrcfs@us.ibm.com or write to:

IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
USA

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. See 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 or Library Server versions of z/OS books in the Internet library at www.ibm.com/systems/z/os/zos/bkserv/.

One exception is command syntax that is published in railroad track format, which is accessible using screen readers with IBM Knowledge Center, as described in “Dotted decimal syntax diagrams.”

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users accessing IBM Knowledge Center using a screen reader. In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always

© Copyright IBM Corp. 2000, 2015

present together (or always absent together), they can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read out punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, you know that your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, it is preceded by the backslash (\) character. The * symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is given the format 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol giving information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, this indicates a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should see separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- A question mark (?) means an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element, (for example 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- An exclamation mark (!) means a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicate that the syntax element is the default option for all syntax elements that share the same dotted

decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP will be applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP applies only to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.

- An asterisk (*) means a syntax element that can be repeated 0 or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data area, you know that you can include one data area, more than one data area, or no data area. If you hear the lines 3*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

  Notes:

  1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
  2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you could write HOST STATE, but you could not write HOST HOST.
  3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.

- + means a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times; that is, it must be included at least once and can be repeated. For example, if you hear the line 6.1+ data area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. Similar to the * symbol, the + symbol can only repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Notices

This information was developed for products and services offered in the USA.

IBM may not offer all of the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

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

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

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

© Copyright IBM Corp. 2000, 2015

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 want 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.

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.

If you are viewing this information softcopy, the photographs and color illustrations might not appear.

COPYRIGHT LICENSE:

This information might contain 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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_.

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.

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.

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 acknowledgment:

```
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. 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.

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.

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 acknowledgment:

   `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.

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)

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 acknowledgment: "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 acknowledgment:

   "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.

The IBM implementation of the Resource Reservation Setup Protocol API (RAPI) described in this document is derived from the Resource Reservation Setup Protocol API (RAPI) specification, The Open Group Document Number C809, ISBN 185912264, published by The Open Group, December 1998.

The specification document is copyrighted by The Open Group. See The Open Group website at http://www.opengroup.org/publications/catalog/c809.htm for

the source of the specification. See *z/OS Communications Server: IP Programmer's Guide and Reference* for details about IBM use of this information.

Copyright © 2004 IBM Corporation and its licensors, including Sendmail, Inc., and the Regents of the University of California.

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

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, the photographs and color illustrations may not appear.

You can obtain softcopy from the z/OS Collection (SK3T-4269), which contains BookManager and PDF formats.

Minimum supported hardware

The minimum supported hardware for z/OS releases identified in z/OS announcements can subsequently change when service for particular servers or

devices is withdrawn. Likewise, the levels of other software products supported on a particular release of z/OS are subject to the service support lifecycle of those products. Therefore, z/OS and its product publications (for example, panels, samples, messages, and product documentation) can include references to hardware and software that is no longer supported.

- For information about software support lifecycle, see: http://www-01.ibm.com/software/support/systemsz/lifecycle/
- For information about currently-supported IBM hardware, contact your IBM representative.

Programming interface information

This publication documents information NOT intended to be used as Programming Interfaces of z/OS Communications Server.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, HCD, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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

Other product and service names might be trademarks of IBM or other companies.

Bibliography

This bibliography contains descriptions of the documents in the z/OS Communications Server library.

z/OS Communications Server documentation is available in the following forms:

- Online at the z/OS Internet Library web page at www.ibm.com/systems/z/os/zos/bkserv/
- In softcopy on CD-ROM collections. See "Softcopy information" on page xxix.

z/OS Communications Server library updates

An index to z/OS Communications Server book updates is at http://www.ibm.com/support/docview.wss?uid=swg21178966. Updates to documents are also available on RETAIN and in information APARs (info APARs). Go to http://www.ibm.com/software/network/commserver/zos/support to view information APARs.

z/OS Communications Server information

z/OS Communications Server product information is grouped by task in the following tables.

Planning

Title	Number	Description
z/OS Communications Server: New Function Summary	GC27-3664	This document is intended to help you plan for new IP or 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.
z/OS Communications Server: IPv6 Network and Application Design Guide	SC27-3663	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
z/OS Communications Server: IP Configuration Guide	SC27-3650	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 with the z/OS Communications Server: IP Configuration Reference.

© Copyright IBM Corp. 2000, 2015

Title	Number	Description
z/OS Communications Server: IP Configuration Reference	SC27-3651	This document presents information for people who want to administer and maintain IP. Use this document with the z/OS Communications Server: IP Configuration Guide. The information in this document includes: • TCP/IP configuration data sets • Configuration statements • Translation tables • Protocol number and port assignments
z/OS Communications Server: SNA Network Implementation Guide	SC27-3672	This document presents the major concepts involved in implementing an SNA network. Use this document with the z/OS Communications Server: SNA Resource Definition Reference.
z/OS Communications Server: SNA Resource Definition Reference	SC27-3675	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 with the z/OS Communications Server: SNA Network Implementation Guide.
z/OS Communications Server: SNA Resource Definition Samples	SC27-3676	This document contains sample definitions to help you implement SNA functions in your networks, and includes sample major node definitions.
z/OS Communications Server: IP Network Print Facility	SC27-3658	This document is for systems 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.

Operation

Title	Number	Description
z/OS Communications Server: IP User's Guide and Commands	SC27-3662	This document describes how to use TCP/IP applications. It contains requests with which a user can 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.
z/OS Communications Server: IP System Administrator's Commands	SC27-3661	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.
z/OS Communications Server: SNA Operation	SC27-3673	This document serves as a reference for programmers and operators requiring detailed information about specific operator commands.
z/OS Communications Server: Quick Reference	SC27-3665	This document contains essential information about SNA and IP commands.

Customization

Title	Number	Description
z/OS Communications Server: SNA Customization	SC27-3666	This document enables you to customize SNA, and includes the following information: • 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
z/OS Communications Server: IP Sockets Application Programming Interface Guide and Reference	SC27-3660	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.
z/OS Communications Server: IP CICS Sockets Guide	SC27-3649	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.
z/OS Communications Server: IP IMS Sockets Guide	SC27-3653	This document is for programmers who want application programs that use the IMS TCP/IP application development services provided by the TCP/IP Services of IBM.
z/OS Communications Server: IP Programmer's Guide and Reference	SC27-3659	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.
z/OS Communications Server: SNA Programming	SC27-3674	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.
z/OS Communications Server: SNA Programmer's LU 6.2 Guide	SC27-3669	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.)
z/OS Communications Server: SNA Programmer's LU 6.2 Reference	SC27-3670	This document provides reference material for the SNA LU 6.2 programming interface for host application programs.
z/OS Communications Server: CSM Guide	SC27-3647	This document describes how applications use the communications storage manager.

Title	Number	Description
z/OS Communications Server: CMIP Services and Topology Agent Guide	SC27-3646	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
z/OS Communications Server: IP Diagnosis Guide	GC27-3652	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.
z/OS Communications Server: ACF/TAP Trace Analysis Handbook	GC27-3645	This document explains how to gather the trace data that is collected and stored in the host processor. It also explains how to use the Advanced Communications Function/Trace Analysis Program (ACF/TAP) service aid to produce reports for analyzing the trace data information.
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	GC27-3667 GC27-3668	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.
z/OS Communications Server: SNA Data Areas Volume 1 and z/OS Communications Server: SNA Data Areas Volume 2	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
z/OS Communications Server: SNA Messages	SC27-3671	This document describes the ELM, IKT, IST, IUT, IVT, and USS messages. Other information in this document includes: • Command and RU types in SNA messages • Node and ID types in SNA messages • Supplemental message-related information
z/OS Communications Server: IP Messages Volume 1 (EZA)	SC27-3654	This volume contains TCP/IP messages beginning with EZA.
z/OS Communications Server: IP Messages Volume 2 (EZB, EZD)	SC27-3655	This volume contains TCP/IP messages beginning with EZB or EZD.
z/OS Communications Server: IP Messages Volume 3 (EZY)	SC27-3656	This volume contains TCP/IP messages beginning with EZY.
z/OS Communications Server: IP Messages Volume 4 (EZZ, SNM)	SC27-3657	This volume contains TCP/IP messages beginning with EZZ and SNM.
z/OS Communications Server: IP and SNA Codes	SC27-3648	This document describes codes and other information that appear in z/OS Communications Server messages.

Index

Numerics

0.0.0.0 address in traps from the SNMP agent 631
64-bit common storage 58

A

abends 350
  CEEDUMP 21
  dumps of ASIDs 50
  FTP client 491
  FTP server 453
  LPD (line printer daemon) 422
  NCPROUTE 797
  OMPROUTE 762
  platform 21
  SMTP (simple mail transfer protocol) 536
  SNALINK LU0 558
  SNALINK LU6.2 589
  SNMP (simple network management protocol) 614
  SNTP (simple network time protocol) 913
  system 21
  TN3270 Telnet client 524
  tn3270E Telnet server 516
  user 21
access control problems 417
accessibility 1027
active routes versus static routes 961
ADDRBLOK data set 542
addressing on the Internet 961
Aggressive mode 973
allocation tables, DU 189
AMBLIST 16
APAR (authorized program analysis report) 3
API, TCPIPCS subcommand 179
APPE for data set fails, FTP 461
ARP cache, querying 40
AT-TLS
  common errors 712
  common startup errors 707
  diagnosing problems 707
  subcommand 709
  traces 709
authorized program analysis report (APAR) 3
autolog, FTP server problems with 452

B

batch job, formatting component traces 57
BEGINROUTES statement
  configuring routing tables 969
  problems uncovered by NETSTAT 39
  problems with X.25 818, 820, 824
  subnet masks 966
  use with NCPROUTE 799, 800
  use with OMPROUTE 761
bridge 955, 957, 961

C

C_INET PFS
  OMPROUTE behavior in 762
CEEDUMP
  description 21
  for SNMP abends 614
CICS (customer information control system) 855
Class D group addresses 963
commit bit
  Quick mode 978
  Quick mode (phase 2) SA states 980, 981
common INET
  OMPROUTE behavior in 762
common storage tracking 16
Communications Server for z/OS, online information xxx
component ID, TCP/IP 18
component trace
  DISPLAY TRACE command 49
  displaying status 49
  general description 14
  managed data set 8
  OMPROUTE 759
  stopping 50
  TN3270E Telnet server 524
  use with OMPROUTE, see also OMPROUTE application trace 786
  viewing data without an abend 50
component trace, changing options
  with PARMLIB member 46
  without PARMLIB member 48
CONFIG, TCPIPCS subcommand 181
configuration values, displaying device 39
CONNECTION, TCPIPCS subcommand 182
consistency check messages, multilevel security 419
control block addresses, displaying TCP/IP 225
control blocks, CICS sockets 859
control blocks, displaying 267
  MTCB (master TCP control block) 250
  MUCB (master UDP control block) 270
  socket 223
  stream 248
  TCBs (TCP control blocks) 250
  Telnet 252
  timer 254
  UCBs (UDP control blocks) 270
control, access problems 417
COUNTERS, TCPIPCS subcommand 184
CSM storage 943
CTRACE command syntax 68
CTRACE disabling, IKE daemon 354, 355
CTRACE enabling, IKE daemon 354
CTRACE status, IKE daemon 355
CTRACE TRACERES records, disable 910
CTRACE TRACERES records, enable 910
CTRACE TRACERES records, format 911
CTRACE, IKE daemon 354
CTRACE, IPCS subcommand 53
CTRACE, resolver 907
CTRACE, TRACERES 909
customer information control system (CICS) 855
customer number 18

© Copyright IBM Corp. 2000, 2015

D

daemon problems, IKE 334
data block, displaying 303
data buffers, displaying 303
data link control (DLC) 571
data traces
  displaying 147
  general description 14
  starting 146
date and time formatting 306
DATTRACE command 146
DB2
  query support for FTP client 480
  query support for FTP server 447
DBCS (double-byte character set)
  FTP client 475
  FTP server 446
DDNS, location of trace output 10
debug information, IKE 350
debug information, IkeSyslogLevel 351
debug information, PagentSyslogLevel 351
debug traces
  for SNALINK LU0 561
  for z/OS UNIX applications 14
  SNTP (simple network time protocol) 913
debugging switches, z/OS UNIX sendmail 547
DEBUGTRC 777
DELE for data set fails, FTP 461
device interface, connecting 182, 184
device interface, displaying 53, 181
diagnosing FTP client user exit problems 489
diagnosing FTP transfer failures with EZA2589E 483
diagnostic aids for IPv6 support 551
disability 1027
dispatchable unit control block (DUCB), displaying 186
dispatchable units, displaying TCP/IP 225
DISPLAY command, use with LUs 581
display commands 16
DISPLAY TCPIP,,OSAINFO command 40
DISPLAY TCPIP,,STOR 16
distributed protocol interface (DPI) 609, 796
DLC (data link control) 571
DNS, online information xxxii
double-byte character set (DBCS)
  FTP client 475
  FTP server 446
DPI (distributed protocol interface) 609, 796
DU allocation tables, displaying 189
DUAF, TCPIPCS subcommand 186
DUCB, displaying 186
DUCB, TCPIPCS subcommand 189
DUMP command 50
dump, selecting 8
dumps
  ABEND 8
  definition 8, 21
  dumps
    SVC 8
  FFST 13
  stand-alone 13
  SVC 13
  SYSABEND 13
  SYSMDUMP 13
  SYSUDUMP 13
  to obtain component trace data 50
Dynamic VIPA 383

E

enable, disable, and display the status of the OMPROUTE CTRACE 790
EPSV commands fail 457
errno reason code, displaying 276
ERRNO, IPCS subcommand 276
errnojr reason code, displaying 276
error exit codes, FTP server 444
establishing security associations, IKE 334
Ethernet
  defining for NCPROUTE 797
  DIX V2 959
event trace (SYSTCPIP)
  TCP/IP stacks 58
  Telnet 58
execution environments
  FTP client 474
Express Logon
  analyzing start problems with the DCAS 866
  client interface problems 867
  overview 865
extended tracing, FTP 471
EZA2589E, diagnosing FTP transfer failures 483
EZAFTPAP start procedure 444
EZD1215-EZD1234 419

F

FDDI (fiber distributed data interface) 957
FFST (first failure support technology)
  dump 13
fiber distributed data interface (FDDI) 957
field maintenance ID 18
file tagging support
  FTP client 476
file tagging support, FTP client
  ASCII file transfers 476
  binary file transfers 479
file transfer protocol (FTP), see also FTP 443
File transfers, restarting 482
firewall, FTP connections 494
formatting component traces using a batch job 57
formatting event trace records for TCP/IP stacks 68
formatting IKE daemon trace records 355
formatting packet trace using a batch job 145
FRCA, TCPIPCS subcommand 192
FTP client 473
  customizing configuration parameters 474
  DB2 480
  DBCS (double-byte character set) 475
  execution environments 474
  exit 475
  file tagging support 476
  MBCS (mult-byte character set) 476
  naming considerations 475
  naming MVS data sets 445
  naming z/OS UNIX file system files 445
  overview 474
  restarting file transfers 482
  SBCS (single-byte character set) 475
  SQL support 480
  trace destination 498
  translation and data conversion support 475
  using FTP.DATA 474
  using TCPIP.DATA 474
  using traces 497

FTP client problems
"Unknown Host" message 493
abends 491
AT-TLS 456
client rejects the RESTART subcommand 482
client rejects the SRESTART subcommand 482
client user exit problems 489
configuration values 491
connection through SOCKS server to FTP server fails 492
conversion between <codeset> and UCS-2 494
data connection 494
data transfer 494
diagnosing FTP transfer failures with EZA2589E 483
documentation for IBM Support Center 499
firewall 494
logging into the server 492
PDS member statistics 497
server rejects password 493
server rejects PORT or EPRT command with 504 replies 495
SOCKS configuration file 492
SOCKS connection fails 492
starting 491
tracing 497
wrong SOCKS server 492
FTP server
controlling traces with MODIFY operator command 467
customizing server parameters 444
customizing the start procedure 444
DB2 query support 447
DBCS (double-byte character set) 446
definitions and setup 444
error exit codes 444
LOADLIB directory information 463
location of trace output 9
logging activity 450
MBCS (multibyte character set) 447
name considerations 444
naming MVS data sets 445
naming z/OS UNIX file system files 445
overview 443
role of TCPIP.DATA 444
SBCS (single-byte character set) 446
SQL 447
SQL support 447
traces 464
translation and data conversion support 446
FTP server problems
abends 453
AT-TLS 456
autolog 452
checkpoint markers 463
client abends during RETR 462
command failure 461
common z/OS UNIX 451
configuration values 452
connection 454, 456
data set allocation 459, 460
data set disposition 463
data transfer 457
documentation for IBM Support Center 472
incomplete initialization 451
job entry subsystem (JES) 449
login failure 455
lost messages and traces 453
MVS data set 461
password validation fails 454
FTP server problems *(continued)*
port specification 452
session 454
terminated data transfer 462
user exit routine 453
working directory 456
FTP session, IPv4 494
FTP.DATA data set
definition 444
specifying working directory 456
FTP.DATA, TRACE statement in 465

G

gateway (router)
definition 955
indirect routing 965
internet addressing 961
LPR and LPD problems 422
using Ping to check connectivity 34
gateway IP address 220
GATEWAY statement
NCPROUTE GATEWAYS data set 806
general tracing, FTP 467
global work area (GWA) 859
GWA (global work area) 859

H

hangs, collecting documentation 23
hardware problem 7
hash tables, displaying 194
HASH, TCPIPCS subcommand 194
header fields, displaying ICMPP 278
header fields, displaying IP 280
header fields, displaying TCP 305
header, displaying Resolver 282
header, displaying UDP 307
header, system dump 198
HEADER, TCPIPCS subcommand 198
HELP, TCPIPCS subcommand 200

I

IBM Software Support Center, contacting xxiv
IBM Software Support Center, documentation for
FTP client problems 499
FTP server problems 472
LPD (line printer daemon) problems 422
SNMP problems 614, 621
X.25 NPSI 825
IBM Software Support Center, when to call 3
IBM Support Center, documentation for
general information 18
server connection problems 41
ICMPP header fields, displaying 278
ICMPPHDR, IPCS subcommand 278
IEEE 802.3 959
IEEE 802.5, relationship between RC and I fields 958
IKE daemon
abends 350
aggressive mode 973
CTRACE 354, 355
debug information 350
diagnosing problems 333
establishing security associations 334

IKE daemon *(continued)*
formatting trace records 355
initialization problems 334
interpreting IKE daemon phase 1 SA states 974
interpreting IKE daemon phase 2 SA states 980
ISAKMP Main mode limitations 982
Main mode 972
overview 333, 971
phase 1 exchanges 971
phase 2 977
problems 334
Quick mode 977
Quick mode with commit bit 978
specifying options after initialization 355
specifying options at initialization 352
specifying trace options 351
TCP/IP services component trace 351
IkeSyslogLevel 351
IMS socket interface
configuration 829
NETSTAT command with 846
overview 827
IMS socket interface problems
bad connections 833
building a component 841
configuration mistakes 830
data transfer 838
error messages and return codes 837, 847
starting and stopping components 831
unexpected database actions 842
IMS socket interface traces 845, 846
inetd.conf file 599
Information APARs xxviii
initialization problems, FTP daemon 451
initialization problems, IKE 334
interactive problem control system, see also IPCS 16
interface, connecting 182, 184
interface, displaying 53, 181
Internet
addressing 961
protocol (IP) 961
Internet, finding z/OS information online xxx
interpreting IKE daemon phase 2 SA states 980
IP header fields, displaying 280
IP security
diagnosing problems, overview 733
problems 733
IPCS (interactive problem control system)
definition 16
sending print output 303
IPCS panels 54
IPCS subcommand
CTRACE 53
ERRNO 276
RESOLVER 282
SETPRINT 303
SKMSG 303
TCPHDR 305
TCPIPCS 176
TOD 306
UDPHDR 307
IPHDR, IPCS subcommand 280
IPSEC, TCPIPCS subcommand 201
IPv4 FTP session 494
IPv4 subnetting 966
IPv6 support, diagnostic aids 551
ISAKMP Main mode limitations 982

J

JES (job entry subsystem), output not found 449
job entry subsystem (JES), output not found 449

K

keyboard 1027

L

Language Environment run-time library 18
leakage problems, memory allocation 682
license, patent, and copyright information 1031
Limiting CSM storage utilization 943
line printer daemon, see also LPD 421
line printer requester, see also LPR 421
LIST MODIFY command 580
Load Balancing Advisor
abends 316
Advisor and Agent problems 316
Advisor or Agent appears to be hung 318
debug settings and corresponding syslogd priority levels 320
diagnostic data 315
group names in displays are indecipherable 319
overview 315
syslogd priority levels 320
workload distribution 316, 317
load module transfer failures 458
load modules 8, 16
LOCK, TCPIPCS subcommand 204
locks, displaying 204
LOCKSUM option 204
logon problems, TN3270E Telnet server 516
LOGREC data set 16
loops
collecting documentation 22
definition 22
SMTP 540
LPD (line printer daemon)
abends 422
activating server traces 431
DEBUG option 431
overview 421
timeouts, hangs, waits 422
unsupported filter 431
LPR (line printer requester)
activating client traces 425
creating client traces 425
FILTER X option 425
garbled data 423
identifying a port 425
missing data 423
nonworking options 424
overview 421
timeouts, hangs, waits 422
XLATE option 425

M

Main mode 972
scenarios 982, 983
mainframe
education xxviii
management information base (MIB) 609, 796

MAP, TCPIPCS subcommand 205
master TCB control block, displaying 250
master trace 14
master UDP control block (MUCB), displaying 270
maximum transmission unit (MTU) 956
MAXSOC value 90
MBCS (multibyte character set)
  FTP client 476
  FTP server 447
memory allocation, leakage problems 682
message block, displaying the stream 303
message EZYFT47I 491
message suppression 453
messages, EZD1215-EZD1234 419
MIB (management information base) 609, 796
MODIFY command
  controlling the FTP server traces 467
  with NCPROUTE 806
module ID, displaying 276
module tables, displaying 208
motif 605
MTABLE, TCPIPCS subcommand 208
MTU (maximum transmission unit) 956
MUCB (master UDP control block), displaying 270
multibyte character set (MBCS)
  FTP client 476
  FTP server 447
multicast, Class D addresses 963
multilevel security consistency check messages 419
multiple stacks 26, 52, 615, 762
MVS data set
  FTP naming conventions 445
  FTP server unable to find 461
MVS system console log for REXECD 593

N

name considerations for z/OS UNIX FTP 444
NAT firewall 494
NCPROUTE
  abends 797
  activating global traces 806
  activating selective traces 806
  bad connections 797
  communication with SNMP 796, 797
  defining 796
  incorrect output 803
  overview 793
  PING failure 801
  RIP (routing information protocol) commands 795
  session outages 804
  trace example 807
negotiating security associations 971
netdata, z/OS UNIX Telnet 502
Netstat ARP/-R 40
netstat command 39
  displaying ADDRESS and LINK values 39
  displaying device configuration values 39
  displaying routing tables 39
  Netstat DEVLINKS/-d 39
  Netstat HOME/-h 39
  Netstat ROUTE/-r 39
  querying ARP cache 40
NETSTAT command
  diagnosing a timeout 38
  displaying data trace 147
  verifying SNALINK LU6.2 580
NETSTAT command *(continued)*
  with IMS socket interface 845, 846
  with SNALINK LU0 559
Netstat DEVLINKS/-d 39
Netstat HOME/-h 39
Netstat ROUTE/-r 39
NetView 610, 611, 620, 803
network IP addresses 963
Network SLAPM2 subagent
  statements 612
network, problems with SNA 574, 576

O

OMPROUTE
  abends 762
  client cannot reach destination 762
  component trace support for 786
  connection problems 762
  overview 759
OMPROUTE CTRACE 790
OMPROUTE subagent
  start options 612
  statements 612
onetstat command
  diagnosing a timeout 38
  OMPROUTE problem diagnosis 762
  use with data trace 147
  viewing host addresses 26
open shortest path first (OSPF) 759
operating system identification 18
OPTIONS keywords 101
orexec (remote execution protocol), z/OS UNIX
  trace example 600
orexecd (remote execution protocol daemon), z/OS UNIX,
  activating debug trace 602
orshd (remote shell daemon), z/OS UNIX
  activating debug trace 603
OSPF (open shortest path first) 759
otelnetd, diagnostic messages sent to wrong file 501

P

packet trace
  general description 14, 94
  starting 95
  system trace 14
  traces
    system trace 14
  with IMS socket interface 845
  with SNALINK LU0 561
  with SNALINK LU6.2 579, 587
PagentSyslogLevel 351
PASV and EPSV commands fail 457
PASV commands fail 457
patricia trees, displaying 260
PDUs (protocol data units)
  NCPROUTE 794
  SNMP (simple network management protocol) 609
PFS (physical file structure) 146
phase 1 971
phase 1 SA states 974
phase 2 977
physical file structure (PFS) 146
Ping command
  SNMP (simple network management protocol) 41

Ping command *(continued)*
using 34
verifying device connection 34
verifying network definition 34
verifying onetstat -h 34
verifying packets sent and received 34
verifying route back to local host 34
verifying route definition 34
verifying router can forward 34
verifying TCP/IP 34
PING command
NCPROUTE client problems 801
with IMS socket interface 831
X.25 NPSI 825
Ping timeout 38
PKTTRACE statement 94
platform abends 21
policies 687
policy agent problems
gathering diagnostic information 659
initialization problems 662
LDAP problems 678
leakage problems 682
memory allocation 682
overview 657
policy definition problems 663
service policy 658
service policy scope 658
policy agent terms
differentiated services 658
integrated services 658
quality of service (QoS) 658
resource reservation protocol (RSVP) 658
service differentiation 658
service level agreement (SLA) 658
service policy 658
Policy Agent, location of trace output 11
POLICY, TCPIPCS subcommand 209
popper, diagnostic aids for z/OS UNIX 554
Popper, location of trace output 10
prerequisite information xxviii
print output, sending IPCS 303
problem number 18
PROFILE, TCPIPCS subcommand 211
program objects 16
protocol data units (PDUs)
NCPROUTE 794
SNMP (simple network management protocol) 609
PROTOCOL, TCPIPCS subcommand 215
ptydata, z/OS UNIX Telnet 502

Q

QoS
policy 658
policy scope 658
Quick mode
(phase 2) SA states with commit-bit support 981
(phase 2) SA states without commit-bit support 980
overview 977
with commit bit 978

R

RACF (resource access control facility), VARY command and 146
RAW, TCPIPCS subcommand 218
reason codes, displaying 276
register save area (RSA) 186
rejected password, FTP 493
release number, TCP/IP 18
remote execution protocol daemon, see also REXECD 591, 599
remote execution protocol, see also REXEC 591, 599
remote shell client, see also RSH 591
remote shell daemon, see also RSH 599
Resolver 869
Resolver header, displaying 282
RESOLVER trace 544
RESOLVER, IPCS subcommand 282
resource access control facility (RACF), VARY command and 146
resource measurement facility 16
RESTART subcommand, client rejects 482
RETAIN database 3
RETR for data set fails, FTP 461
REXEC (remote execution protocol), non-z/OS UNIX
activating debug trace 592
overview 591
problem documentation 591
trace example 592
REXEC (remote execution protocol), z/OS UNIX
debug trace 600
trace example 600
REXECD (remote execution protocol daemon), non-z/OS UNIX
command options for tracing 593
overview 592
problem documentation 593
trace example 593
REXECD (remote execution protocol daemon), z/OS UNIX
trace example 602
REXECD (remote execution protocol daemon), z/OS UNIX, activating debug trace 602
REXX executables for TCPIPCS 176
RFC (request for comments) 1003
accessing online xxx
RIP (routing information protocol)
NCPROUTE implementation 794, 803
OMPROUTE implementation 759
RNFR for data set fails, FTP 461
RNTO for data set fails, FTP 461
ROUTE, TCPIPCS subcommand 220
router (gateway)
definition 955
indirect routing 965
internet addressing 961
LPR and LPD problems 422
using Ping to check connectivity 34
routing
datagram algorithm 965
datagram algorithm with subnets 967
direct 964
dynamic 969
dynamic IPv4 970
dynamic IPv6 970
indirect 965
static 969
routing control blocks, displaying 220
routing information protocol (RIP)
NCPROUTE implementation 794, 803
OMPROUTE implementation 759
routing information, displaying 39

routing table entry, displaying 220
RSA (register save area) 186
RSH 601
RSH (remote shell client), non-z/OS UNIX
  client trace using SEND command 595
  overview 591
  trace example 592
RSH, z/OS UNIX
  debug trace 601
  trace example 601
RSHD (remote shell daemon), z/OS UNIX
  activating debug trace 603
  random errors 603
  trace example 603
RSVP agent problems
  application problems 689
  gathering diagnostic information 688
  initialization problems 688
  log file example 690
  overview 685
  reservation objects 686
  reservation styles 686
  reservation types 686
  service policy problems 689
RSVP Agent, location of trace output 11
RSVP processing 687

S

SBCS (single-byte character set)
  FTP client 475
  FTP server 446
secure IPv4 FTP session 494
secure sockets layer (SSL) encryption 523
security associations
  aggressive mode 973
  common problems 335
  interpreting IKE daemon phase 1 SA states 974
  interpreting IKE daemon phase 2 SA states 980
  ISAKMP Main mode limitations 982
  Main mode 972
  overview of negotiating 971
  phase 1 971
  phase 2 977
  problems establishing 334
  Quick mode 977
  Quick mode with commit bit 978
selecting a dump 8
selecting a service aid 12
selecting a trace 8
sendmail, diagnostic aids for z/OS UNIX 547, 549
server data block, displaying the TSDB 263
server data extension, displaying the TSDX 265
service aid, selecting 12
service aid, selection 7
service level indication processing (SLIP) 8, 16
service policy scopes
  DataTraffic 658
  RSVP 658
services component trace, TCP/IP 351
session problems
  FTP 454
  SNALINK LU0 hangs 558
  SNALINK LU0 outages 560
  tn3270 Telnet client hangs 525
  TN3270E Telnet server hangs 518
  TN3270E Telnet server outages 521
session problems *(continued)*
  X.25 NPSI hangs 824
SETPRINT, IPCS subcommand 303
shortcut keys 1027
simple mail transfer protocol, see also SMTP 535
simple network management protocol, see also SNMP 609
simple network time protocol, see also SNTP 913
single-byte character set (SBCS)
  FTP client 475
  FTP server 446
SKMSG, IPCS subcommand 303
SLIP (service level indication processing) 8, 16
SMTP (simple mail transfer protocol)
  abends 536
  ADDRBLOK data set 542
  defining 536
  delivery problems 537
  environment 535
  incorrect output 540
  looping problems 540
  re-resolution of queued mail 541
  receiver overview 535
  RESOLVER trace 544
  sender overview 535
  spooling problems 537
SMTP, location of trace output 9
SNALINK LU0
  abends 558
  location of trace output 10
  session hangs 558
  session outages 560
SNALINK LU6.2
  address space problems 569, 580
  configuration mistakes 567
  data corruption problems 579
  data loss problems 577
  DLC connection problems 571
  location of trace output 10
  network components 565, 566
  network problems 574, 576
  traces 584, 588, 589
  using NETSTAT 580
SNAP (subnetwork access protocol) 960
SNMP (simple network management protocol)
  abends 614
  agent 614
  agent does not respond 627
  client overview 610
  communication with NCPROUTE 796
  connection problems 615
  I/O error using PING 631
  incorrect output 621
  interaction with NCPROUTE 793
  MIB 609
  multiple stack problems 615
  NetView 610
  query engine 620, 621, 651
  snmp command 610, 611, 633
  snmpd command 610
  socket calls 613
  starting with snmp 633
  subagent 612
  traces 633, 634, 635
  unknown variable in output 622
  variable format incorrect 624
  variable value incorrect 626

SNMP agent
configuration data sets 611
SNMP agent, 0.0.0.0 address in traps 631
snmp command 610, 611, 633
SNMP problems
related to NCPROUTE 798
SNMPIUCV subtask 620
SNMP traces
agent trace example 639
query engine IUCV trace example 651
query engine trace example 640
subagent trace example 639
SNMPIUCV subtask 620
SNMPv3
AES encryption or decryption failures 629
SNTP (simple network time protocol) 913
abends 913
debug trace 913
socket calls, SNMP (simple network management protocol) 613
socket control block, displaying 223
socket data, tracing 146
SOCKET, TCPIPCS subcommand 223
softcopy information xxviii
specifying options after initialization 355
specifying options at initialization 352
specifying trace options 351
SQL
FTP client support 480
FTP server support 447
problems with FTP 448, 480
SRESTART subcommand, client rejects 482
SSA REXX/370 runtime libraries 176
SSL (secure sockets layer) encryption 523
stand-alone dumps 13
start procedure, FTP server 444
starting TN3270E Telnet subagent trace 636
state, displaying TCP/IP 225
STATE, TCPIPCS subcommand 225
static routes compared to active routes 961
status, displaying device 39
STDOUT, use by OMPROUTE 759
STOR for data set fails, FTP 461
storage map, displaying 205
storage, displaying TCP/IP 225, 246
STORAGE, TCPIPCS subcommand 246
stream control blocks, displaying 248
stream message block, displaying 303
STREAM, TCPIPCS subcommand 248
subagent programs, user-written for SNMP (simple network management protocol) 609
subagent, Network SLAPM2 612
subagent, OMPROUTE 612
subagent, SNMP (simple network management protocol) 612
subagents 611
subnetwork access protocol (SNAP) 960
subtasks, displaying TCP/IP 225
summary of changes xxxiii, xxxiv
SVC dumps 13
SWSA 403
syntax diagram, how to read xxv
syntax for the CTRACE command 68
SYSABEND dump 13
SYSERROR data set 3
syslogd
SNMP (simple network management protocol) agent
traces 614, 615
syslogd *(continued)*
traces for FTP server 465
traces for OMPROUTE 759
traces for REXECD 602
traces for RSHD 603
using with z/OS UNIX Telnet 501
SYSMDUMP 13
sysplex
routing problems 409
SWSA 403
Sysplex distributor 383
diagnosing 384
overview 383
sysplex-wide dynamic source VIPAs 396
Sysplex-wide Security Association (SWSA) 403
SYSPLEXPORTS problems 397
SYSPRINT data set 3
SYSTCPIS 149, 701
system abends 21
system dump header, displaying 198
SYSUDUMP 13

T

task interface element (TIE) 859
task related user exit (TRUE), CICS sockets 860
tasks
(gerund phrase)
step for 601
(noun, gerund phrase)
step 949, 950
steps 946
access control support
overview 418
activating LPR client traces
step for 425
activating server traces
step for 431
activating the z/OS UNIX RSHD debug trace
step for 603
analyzing hangs
steps for 23
analyzing incorrect mail output
steps for 541
analyzing incorrect output (client)
steps for 527
analyzing incorrect output (server)
steps for 520
analyzing logon problems (server)
steps for 517
analyzing mail delivery problems
steps for 538
analyzing routing failures
steps for 802
analyzing session hangs
steps for 559
analyzing session hangs (client)
steps for 525
analyzing session hangs (server)
steps for 519
analyzing session outages (server)
steps for 522
checking DLC connection status
steps for 571
checking network connection problems
steps for 574

tasks *(continued)*
collecting documentation
steps for 22
creating client trace output
step for 425
creating server trace output
step for 431
data on the TCP queues, reviewing
steps 947
determining how your IDS policies have been mapped by the stack
step for 700
determining which IDS policies are active in Policy Agent
step for 699
determining why IDS SyslogD output is missing
steps for 700
diagnosing
overview 3
diagnosing a AT-TLS problem
steps for 708
diagnosing child-server transactions not starting
steps for 858
diagnosing CICS listener not initialized
steps for 857
diagnosing CICS socket interface not initialized
steps for 856
diagnosing garbled data
steps for 423
diagnosing hung CICS tasks
steps for 858
diagnosing IKE daemon problems
overview 333
diagnosing incorrect output
steps for 804
diagnosing IP security problems
overview 733
steps for 734
diagnosing LDAP object retrieval problems
steps for 678
diagnosing logon problems
steps for 823
diagnosing LPR working with some options only
steps for 424
diagnosing no response to commands
steps for 849
diagnosing Policy Agent/Sysplex distribution problems
steps for 682
diagnosing problems
steps for 3
diagnosing server connection problems
steps for 26
diagnosing session outages
steps for 805
diagnosing SNMP problems
steps for 613
diagnosing SWSA problems
steps for 403
diagnosing sysplex distributor problems
overview 383
diagnosing sysplex problems
steps for 384
diagnosing sysplex routing problems
steps for 409
diagnosing sysplex-wide dynamic source VIPAs for TCP connections problems
steps for 396
tasks *(continued)*
diagnosing SYSPLEXPORTS problems
steps for 397
diagnosing TCP/IP clients unable to connect
steps for 857
diagnosing timeouts, hangs, and waits
steps for 422
diagnosing TRMD problems
steps for 704
diagnosing truncated or missing print data
steps for 423
diagnostic aids for AT-TLS support
steps for 552
disabling the CTRACE after OMPROUTE has started
steps for 791
disabling the CTRACE after the IKE daemon has started
steps for 355
disabling the CTRACE at IKE daemon startup
steps for 354
disabling the CTRACE at OMPROUTE startup
step for 790
displaying the CTRACE status
step for 355, 791
displaying the internal trace
steps for 860
documenting data loss problems
steps for 577
dynamic IPSec protection
steps for 746
enabling the CTRACE after OMPROUTE has started
steps for 790
enabling the CTRACE after the IKE daemon has started
steps for 354
enabling the CTRACE at IKE daemon startup
steps for 354
enabling the CTRACE at OMPROUTE startup
steps for 790
entering a TCP/IP IPCS subcommand
steps for 310
formatting component traces using IPCS panels
steps for 54
invoking z/OS UNIX RSH trace
step for 601
IP security operation
steps for 741
IP security policy enforcement
steps for 749
manual IPSec protection
steps for 744
NCPROUTE connection problems
steps for 798
negotiating security associations
overview 971
network access
steps for 33
obtaining AMATERSE
steps for 18
obtaining component trace data with an external writer
steps for 51
Policy definition problems
step for 663
resolving the hostname
steps for 869
selecting tools and service aids
overview 7
setting up a SNALINK LU6.2 network
steps for 566

tasks *(continued)*
setting up the IMS TCP/IP services socket interface system
steps for 829
starting manager traces
step for 633
starting network SLAPM2 subagent traces
step for 635
starting OMPROUTE subagent traces
step for 635
starting SNMP agent traces
step for 634
starting TCP/IP subagent traces
step for 635
starting Telnet client traces
step for 528
starting TN3270E Telnet subagent traces
step for 636
starting TRAPFWD traces
step for 636
stopping SNTPD
steps for 913
tools and service aids
overview 13
undeliverable mail items
steps for undeliverable mail items 537
using FTP client SQL support
steps for 480
using FTP server SQL support
steps for 447
using the batch option
step for 311
using the CTRACE command
steps for 55
verifying IP routing to a destination
steps for 29
verifying network interface operation
steps for 33
verifying server operation
steps for 27
verifying the configuration
steps for 419
TBEs (trace buffer entries), displaying 256
TCA (trace control area), displaying 256
TCB control blocks, displaying 250
TCB, TCPIPCS subcommand 250
TCBSUM, subcommand for TCIPICS 250
TCP connections problems 396
TCP header fields, displaying 305
TCP/IP
event trace (SYSTCPIP) 58
online information xxx
protocol specifications 1003
TCP/IP services component trace, IKE daemon 351
TCP/IP state, displaying summary 225
TCP/IP storage, displaying summary 246
TCP/IP subagent, statements 612
TCPHDR, IPCS subcommand 305
TCPIP.DATA data set 444
TCPIPCS subcommand
API 179
CONFIG 181
CONNECTION 182
COUNTERS 184
DUAF 186
DUCB 189
FRCA 192
HASH 194
TCPIPCS subcommand *(continued)*
HEADER 198
HELP 200
IPSEC 201
LOCK 204
MAP 205
MTABLE 208
POLICY 209
PROFILE 211
PROTOCOL 215
RAW 218
REXX executables for 176
ROUTE 220
SOCKET 223
STATE 225
STORAGE 246
STREAM 248
symbols 179
syntax 177
TCB 250
TELNET 252
TIMER 254
TRACE 256
TREE 260
TSDB 263
TSDX 265
TTLS 267
UDP 270
VMCF 273
XCF 274
TCPIPCS subcommand, TTLS 267
Technotes xxviii
Telnet
protocol commands and options 532
TELNET, TCPIPCS subcommand 252
Telnet, tn3270
incorrect output from client 527
starting client traces 528
Telnet, TN3270 client
client abends 524
client definitions 524
client overview 524
client session hangs 525
client trace example 528
Telnet, TN3270E Telnet server
associating CTrace entry with client 524
defining server 515
displaying control blocks (TCPIPCS subcommand, Telnet) 252
event trace (SYSTCPIP) 58
incorrect output from server 520
server abends 516
server logon problems 516
server overview 515
server session hangs 518
server session outages 521
SSL encryption 523
Telnet, z/OS UNIX
-t -D all example 502
arrow keys do not respond 501
debug trace options 502
diagnostic messages sent to wrong file 501
improper authority error 501
locked keyboard 501
netdata and ptydata 502
utmp entries 514
termination notification facility (TNF) 849

TIE (task interface element) 859
time and date formatting 306
timeout
  Ping 38
timer control blocks, displaying 254
TIMER, TCPIPCS subcommand 254
TN3270E Telnet subagent 613
TN3270E Telnet subagent trace, starting 636
TNF (termination notification facility) 849
TOD, IPCS subcommand 306
token-ring
  defining for NCPROUTE 797
  IEEE 802.5 958
tool or service aid, selection 7
tool, selection 7
trace buffer entries (TBEs), displaying 256
trace control area (TCA), displaying 256
TRACE CT command 48, 67
trace entries, FTP 453
trace options, specifying 351
trace output, location of
  DDNS 10
  FTP server 9
  Policy Agent 11
  Popper 10
  RSVP Agent 11
  SMTP 9
  SNALINK LU0 10
  TRMD 11
trace records, IKE daemon 355
TRACE RESOLVER
  MVS batch job environment 888
  output location 887
  resolver 884
  TSO environment 887
  z/OS UNIX batch environment 888
  z/OS UNIX shell environment 887
trace, selecting 8
TRACE, TCPIPCS subcommand 256
traceroute command 41
TRACERTE command 41
traces
  AT-TLS 709
  component trace 14
  data trace 146
  data traces 14
  event trace for TCP/IP stacks 58
  for FTP client problems 497
  for FTP server problems 464
  for SNALINK LU6.2 networks 584
  GTF traces 14
  master trace 14
  packet trace 14, 561, 579, 587
  SNALINK LU0 debug 561
  socket API traces 74
  TCP/IP internal 588
  tn3270 Telnet client 528
  TN3270E Telnet server 524
  VTAM buffer 589
  VTAM trace 14
  z/OS UNIX application debug traces 14
trademark information 1039
translation and data conversion support 446
TRAPFWD trace example 655
traps 8
TREE, TCPIPCS subcommand 260
TRMD problems, gathering diagnostic information 704
TRMD, location of trace output 11
TRUE (task related user exist), CICS sockets 860
TSDB, TCPIPCS subcommand 263
TSDX, TCPIPCS subcommand 265
TSO console log from REXEC 592

U

UCBs (UDP control blocks), displaying 270
UDP control blocks (UCBs), displaying 270
UDP header, displaying 307
UDP, TCPIPCS subcommand 270
UDPHDR, IPCS subcommand 307
user abends 21
user exit routine, FTP 453
utmpx file 514

V

VARY TCPIP command 39
VERBEXIT 16
VIPA (virtual IP address) 761
virtual IP address (VIPA) 761
virtual machine communication facility (VMCF) 849
VMCF (virtual machine communication facility) 849
VMCF, TCPIPCS subcommand 273
VTAM buffer trace
  for tn3270 Telnet client session hangs 526
  for TN3270E Telnet server session hangs 519
VTAM trace 14
VTAM, online information xxx

X

X Window System
  trace when XWTRACE=2 605
  trace when XWTRACELC=2 606
X.25 NPSI
  configuring for NPSI 819
  configuring for VTAM 819
  logon problems 823
  overview 817
  session hangs 824
XCF, TCPIPCS subcommand 274

Z

z/OS Basic Skills Information Center xxviii
z/OS UNIX file system
  file naming conventions for FTP 445
z/OS UNIX RSH 601
z/OS, documentation library listing 1041

Communicating your comments to IBM

If you especially like or dislike anything about this document, you can send us comments electronically by using one of the following methods:

Internet email:
comsvrcf@us.ibm.com

World Wide Web:
http://www.ibm.com/systems/z/os/zos/webqs.html

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

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

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.

© Copyright IBM Corp. 2000, 2015

Product Number: 5650-ZOS

Printed in USA

GC27-3652-03