

z/OS Communications Server



SNA Diagnosis Volume 1: Techniques and Procedures

Version 2 Release 1

Note:

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 659.

First edition (September 2013)

This edition applies to version 2, release 1, modification 0 of z/OS (5650-ZOS), and to subsequent releases and modifications until otherwise indicated in new editions.

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

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

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

Fax (USA and Canada):

1+919-254-1258

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

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, telephone number, or FAX 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, 2013.

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

Contents

Figures	ix
--------------------------	-----------

Tables	xv
-------------------------	-----------

About this document	xvii
--------------------------------------	-------------

Who should read this document	xvii
---	------

How this document is organized.	xvii
---	------

How to use this document.	xvii
-----------------------------------	------

Determining whether a publication is current	xvii
--	------

How to contact IBM service	xviii
--------------------------------------	-------

Conventions and terminology that are used in this document	xviii
--	-------

How to read a syntax diagram.	xix
---------------------------------------	-----

Prerequisite and related information	xxi
--	-----

Summary of changes	xxvii
-------------------------------------	--------------

Part 1. Diagnostic techniques	1
--	----------

Chapter 1. Diagnosing VTAM problems: Where to begin	3
--	----------

Determining whether the problem is VTAM or non-VTAM	3
---	---

Common problems and symptoms	4
--	---

Common problems in subarea networks	4
---	---

Descriptions of common problems in subarea networks	8
---	---

Common problems in APPN networks	22
--	----

Descriptions of common problems in APPN networks	26
--	----

Common problems in HPR networks.	38
--	----

Descriptions of common problems in HPR networks	38
---	----

Common symptoms and associated VTAM problem types.	42
--	----

VTAM internal trace (VIT) analysis tool problems	44
--	----

Checklist for isolating the problem	44
---	----

Common symptoms and actions	45
---------------------------------------	----

Documenting an APAR for VIT analysis tool problems	47
--	----

VTAM dump analysis tool problems	47
--	----

Checklist for isolating the problem	47
---	----

Documenting an APAR for dump analysis tool problems	48
---	----

Recommended documentation for VTAM problems	48
---	----

Methods for submitting documentation	53
--	----

Submitting documentation using FTP.	53
---	----

Using electronic transfer through email attachments.	54
--	----

Submitting documentation on tape	54
--	----

Necessary documentation.	55
----------------------------------	----

Chapter 2. Collecting documentation for specific types of problems	57
---	-----------

Common problem determination procedures	57
---	----

Abnormal end (abend)	57
--------------------------------	----

Wait	61
----------------	----

Loop	81
----------------	----

Message problem	87
---------------------------	----

Incorrect output	90
----------------------------	----

Performance problem	94
-------------------------------	----

Storage problem procedure	97
-------------------------------------	----

Documentation problem.	100
--------------------------------	-----

Failing module	100
--------------------------	-----

Symptom string structure	102
------------------------------------	-----

Reporting the problem to IBM.	103
---------------------------------------	-----

Chapter 3. Collecting documentation for TSO/VTAM problems 105

Initial TSO/VTAM problem analysis	105
Logon problems	106
TSO/VTAM abends	109
ABEND0AB	109
ABEND0AC	110
ABEND0AD	110
ABEND15D	110
Parameter initialization problems	110
Hung terminal problems.	111
Incorrect output problems	112
Screen management problems	112
Screen size problems	117
Performance problems	119

Part 2. Diagnostic procedures 121

Chapter 4. Using DISPLAY and MODIFY operator commands 123

Using VTAM DISPLAY commands for problem determination	124
Display buffer pool use	127
DISPLAY CSDUMP	128
DISPLAY CSMUSE	128
DISPLAY EE.	128
DISPLAY EEDIAG.	131
DISPLAY EEDIAG,TEST=YES	136
DISPLAY EEDIAG,TEST=PENDING.	149
Display Enterprise Extender connection network unreachable partner information	149
Display HPR route test	150
Display ID for an RTP connection	152
Display ID for an RTP PU with HPRDIAG=YES.	152
Display ID for an RTP PU with HPRDIAG=YES and CLEAR=ALL	154
Display NCP storage	155
Display path tables	155
Display resource status	156
Display resources in a pending state.	156
Display route status	156
Example: Solving path problems	157
Display route test	163
Display RTPS options	166
Display TDU information	167
Display traces	175
Display VTAM storage	175
Display workload information for a device	175
Using VTAM MODIFY commands for problem determination	177
Issuing the MODIFY CSDUMP command	177
Modifying input/output problem determination.	177
Modifying message module identification	178
Modifying NCP intensive mode recording.	179
Modifying SDLC link level 2 test	179
Issuing the MODIFY TOPO command to clear EE connection network unreachable partner information	180
Modifying tuning statistics	181
Issuing the MODIFY VTAMOPTS command to change start option values	181

Chapter 5. Using dumps. 183

Dumps on MVS operating system	183
Abend dump	183
Coupling facility structures dump	184
FFST dump	184

Stand-alone dump	184
SVC dump	184
Network control program (NCP) dump	185
When to use the NCP dump	185
Using the NCP option on MODIFY DUMP	186
Using the independent NCP dump utility (channel-attached controller only)	187
Communication scanner processor dump (3720, 3725, and 3745 only).	187
Maintenance and operator subsystem dump (3720, 3725, and 3745 only).	188
Formatting and printing dump output	188
IPCS service aids	188
ABDUMP service aid	189
SADMP service aid	190

Chapter 6. Using VTAM dump analysis tools 191

Enhanced VTAM dump analysis tools	191
Operating environment	191
Using VTAM interactive problem control system (IPCS) CLISTS	192
Obtaining online help for CLISTS	192
Debugging CLIST errors.	192
Printing CLIST output	193
Sample VTAM dump analysis functions	193
Using the panel interface	193
Using the IPCS command line.	194
Using the batch option	195
VTAM formatted dump procedures	196
ALL	196
APPLCONV.	197
APPLMODE.	198
APPMODAL	200
APPNBASE	202
ATMDATA	203
CSMALL	205
CSMBUF	208
CSMCMPID	210
CSMOWNER	212
CSMPOOL	214
FINDDSIB	216
FINDSIB	218
FNDADJCP	222
FNDANDCB	223
FNDCOS	224
FNDDECB	225
FNDENDEL.	226
FNDLCB	227
FNDNDREC	228
FNDNDWGT	230
FNDNODE	231
FNDREREC	232
FNDSCCB	233
FNDSITCB	234
FNDTGREC	235
FNDTGWGT	238
HOST	240
HPRIP.	241
ISTVABND	242
ISTVDUMP	244
SDUMP parameter list	246
ISTVMAP	247
ISTVSAVE	249
ISTVSLIP.	251
MNPSC, MNPSD, MNPSF	254

PABSCAN	257
PARTNRLU	261
RDTCHECK	262
RDTFULL	263
RDTHIER	264
RDTSUM	265
ROUTES	266
RTPINFO	267
SES	271
SIBCHECK	275
SPANC	277
SRTFIND	281
STORAGE	282
TOPOLOGY	283
TRSTRACE	285
VITAL	286
VTAM	287
VTBASIC	289
VTBUF	290
VTCVTPAB	291
VTFNDMOD	294
VTMODS	295
VTNODE	297
VTREADYQ	298
VTRPH	299
VTVIT	300
VTVRBLK	302
VTWRE	304

Chapter 7. Using traces 307

Traces provided by VTAM	307
Activating network traces	308
Starting the generalized trace facility (GTF)	322
Formatting and printing trace records	323
Trace output	325
VTAM trace record formats	325
Buffer contents trace	328
Configuration services XID exit (CSX) buffer trace	335
Directory services session management exit (DSME) buffer trace	336
I/O trace	340
QDIOSYNC trace	341
Resource state trace	345
Session management exit (SME) buffer trace	346
SMS (buffer use) trace	348
TGET/TPUT trace for TSO/VTAM	350
Traces provided by NCP	353
Generalized PIU trace	353
Line trace	354
Network controller line trace (3710 only)	358
Scanner interface trace (3720, 3725, and 3745 only)	359
Transmission group trace	359
Traces provided by TCP/IP	362

Chapter 8. Using the VIT analysis tool. 365

Setting up and running the VIT analysis tool	365
Step 1. Record a VIT	366
Step 2. Set up to run the tool	366
Step 3. Create the parameters for the job	368
Step 4. Run the job	369
Step 5. Check the output	369

Return codes	369
Environment	369
Analyzing storage	369
Counting request/response units (RUs).	377
Extracting information from the VIT.	385
Using the timing options	394
Start and stop time	395
Report interval	396
Parameter syntax	396
Using the I/O options	397
Trace wrapped	397
Format output	398
Parameter syntax	398
Operands.	398
Creating your own parameter data set	399
Chapter 9. Using other problem-solving tools	401
Alert messages from NCP	401
Recording NMVT alerts in LOGREC	401
Messages issued for 3745 bus switching	402
Message IST8811	402
Message IST8821	402
Hardware error recording	403
Logical unit connection test (IBMTEST).	403
NCP error recording	404
Patch areas	404
VTAM patch area	404
TSO/VTAM patch area	405
VTAM load module list	405
Using save-area module linkage conventions—Subarea	405
Using save-area module linkage conventions—APPN	407
Part 3. Appendixes	411
Appendix A. Channel programs	413
Channel programs for channel-attached type 2 and type 4 physical units	413
Channel commands for channel-attached type 2 and type 4 physical units	414
Channel program sequences	417
PUNS-related channel programs	418
Channel programs for channel-to-channel adapters (CTCA), multipath channel (MPC), and APPN host-to-host channels	420
Channel programs for activating the CTC connection	420
Channel commands for channel-to-channel (CTC) adapters	422
XID channel program (X-side).	424
XID channel program (Y-side).	425
Channel programs for CTC data transfer: Blocking protocol	426
Channel programs for channel-to-channel adapters: Nonblocking protocol	429
Channel programs for multipath channel (MPC).	434
Channel programs for activating the MPC connection	434
Channel commands for activating the MPC connection	436
Channel programs for MPC data transfer	436
Channel program for HPDT MPC data transfer	438
Channel programs for channel-attached non-SNA 3270 devices.	440
Channel command words	441
Channel programs.	442
Appendix B. Network flows	447
Generic BIND (GBIND) AMRUs	451
Index of generic BIND (GBIND) AMRU flows	451
Resource activation flows	453

Index of resource activation flows	453
Session establishment flows	466
Index of session establishment flows	466
Deactivation and session termination flows	492
Index of deactivation and session termination flows	492
Error detection and recovery and SSCP management services	512
Index of error detection and recovery and SSCP management services flows	512
Appendix C. APPN flows	521
CP-CP session flows	525
Index of CP-CP session flows	525
Directory services flows	533
Index of directory services flows	533
Register resource flows	533
Locate resource flows.	535
LU-LU session flows	556
Index of LU-LU session flows	556
Dependent LU server flows	587
Index of dependent LU server flows.	587
Single subnetwork flows	588
Cross subnetwork flows	605
High-Performance Routing flows	614
Index of High-Performance Routing flows.	614
Appendix D. Control point/control block (CPCB) operation codes	623
Appendix E. Storage and control block ID codes	637
VTAM control block ID codes	637
Appendix F. Installing dump analysis and VIT analysis tools	641
Concatenating target data sets used in the installation.	641
Customizing IPCS interface.	641
Verifying dump formatter panels.	643
Customizing ISPF interface.	644
Verifying trace formatter panels	646
Appendix G. Problem topics in other libraries	649
Appendix H. Architectural specifications.	653
Appendix I. Accessibility	655
Notices	659
Programming interface information	667
Policy for unsupported hardware.	667
Trademarks	667
Bibliography.	669
Index	673
Communicating your comments to IBM	683

Figures

1. Overview of the abend procedure	58
2. Overview of the wait procedure.	62
3. PAB locations	68
4. Normal PABs, extended PABs, and slightly extended PABs.	68
5. Very extended PAB	69
6. Finding LQAB groups	72
7. Finding waiting request elements for an LQAB group	73
8. Pointers to VTAM locks	79
9. Overview of the loop procedure (part 1 of 2)	83
10. Overview of the loop procedure (part 2 of 2)	84
11. Overview of the message procedure	88
12. Overview of the incorrect output procedure.	91
13. Overview of the performance procedure	95
14. Overview of the documentation procedure.	100
15. Overview of the failing module procedure	101
16. Example of a symptom text string in output	102
17. Enterprise Extender inactivity flows	134
18. Enterprise Extender configuration with firewalls.	137
19. Connectivity test without policy-based routing enabled	138
20. EE connectivity test with a single policy-based routing rule enabled	139
21. Basic EE connectivity test (successful connection)	140
22. Basic EE connectivity test (unsuccessful connection).	141
23. Enterprise Extender with multipath routing enabled	143
24. EE connectivity test with multipath routing enabled (part 1 of 2)	144
25. EE connectivity test with multipath routing enabled (part 2 of 2)	145
26. Enterprise Extender with policy-based routing	146
27. EE connectivity test with multiple policy-based routing rules enabled.	148
28. Example of DISPLAY route status output	157
29. Path problem example network configuration.	158
30. Output of a successful route test	164
31. Output of a failed route test	164
32. Route test failure (TG inactive or migration)	165
33. Route test failure (ER not reversible, exceeds maximum length, or not defined)	166
34. One example of how a TDU war might occur.	169
35. Network traces provided by VTAM	308
36. Starting GTF for the VTAM buffer contents trace.	322
37. Format for buffer contents trace records for VTAM API and TSC	326
38. Format for buffer contents trace records for CMIP services API	327
39. Format for line trace records	328
40. Trace points for the user buffer trace for CMIP services	331
41. Example of partial buffer contents trace output	333
42. Example of full buffer contents trace output	333
43. Example of full buffer contents trace output for CMIP services API	334
44. Example of configuration services XID exit (CSX) buffer trace output	336
45. Example of directory services session management exit (DSME) buffer trace output	339
46. Example of session management exit (SME) buffer trace output.	348
47. Example of SMS trace output	349
48. Example of TGET/TPUT trace output	352
49. Example of line trace output (CS type 2)	356
50. Example of line trace output (CS type 3)	357
51. Example of transmission group trace output	361
52. Sample TCP/IP trace of EE data (part 1 of 2)	362
53. Sample TCP/IP trace of EE data (part 2 of 2)	363
54. Sample JCL for VIT analysis	366
55. Sample VIT analysis tool interactive routine	367

56. VTAM internal trace analysis option panel	368
57. VTAM storage analysis option panel	370
58. Storage analysis with a stop time	377
59. VTAM request/response unit counting	378
60. RU code options	381
61. RU modify options	381
62. VIT extraction Boolean expression panel	386
63. VIT extraction template	387
64. Example of VIT extraction	393
65. VTAM timing options panel	395
66. VTAM I/O options panel	397
67. Parameters coded on multiple lines	399
68. Save-area module linkage conventions—subarea	407
69. Save-area module linkage conventions—APPN	409
70. Data areas used by channel programs for PU types 2 and 4	414
71. Format of Write CCWs with chained data	417
72. Example of an XID exchange	421
73. Data areas used for XID channel programs (X-side)	425
74. Data areas used for XID channel programs (Y-side)	426
75. Buffers used for normal data transfer	427
76. Data areas used for normal data transfer (X-side)	428
77. Data areas used for normal data transfer (Y-side)	429
78. Data areas used for normal data transfer (X-side) nonblocking	431
79. Data areas used for normal data transfer (Y-side) nonblocking	433
80. MPC activation flow	435
81. MPC transmit buffers used for normal data transfer	437
82. Indirect address word structure of multipath channel programs for normal data flow	438
83. Basic read Seldom Ending Channel Program structure of HPDT multipath channel	439
84. Example of Seldom Ending Channel Program structure of HPDT multipath channel	439
85. Data areas used by a channel program for channel-attached non-SNA devices	441
86. Sending an ACTLU request for a logical unit (LU)	452
87. Sending an ACTPU request for a communication controller or physical unit (PU)	452
88. Sending a BIND request to a secondary logical unit (SLU)	453
89. Activating a virtual route (VR) and the associated explicit route (ER)	453
90. Activating a channel-attached communication controller	455
91. Activating a link-attached communication controller	456
92. Activating a link (ACTLINK)	456
93. Activating a switched link with takeover	457
94. Activating a cross-subarea link station	457
95. Establishing a switched connection	458
96. Activating a physical unit type 2.0	459
97. Activating a physical unit type 2.0 with load required	459
98. Moving a SYSGENed physical unit	460
99. Moving a dynamically added physical unit	461
100. SSCP takeover of peripheral node logical units	461
101. Activating a logical unit	462
102. Activating an application program and processing an OPEN ACB request	462
103. Activating CDRM with ERP response	463
104. Activating CDRM with COLD response.	463
105. Activating a CDRM with a virtual-route-based transmission group	463
106. Back-to-back gateway NCP request sessions	464
107. Gateway VTAM requests session	465
108. Non-gateway VTAM requests session	465
109. Default partitioning of gateway VTAM responsibility spanning three networks	468
110. Multiple gateway VTAMs and back-to-back gateway NCPs	469
111. Primary logical unit initiate, OPNDST ACQUIRE	470
112. Primary logical unit initiate, SIMLOGON	471
113. Primary logical unit initiate, SIMLOGON(RELREQ).	472
114. Primary logical unit initiate, SIMLOGON(RELREQ): Session is pending active or already in progress	473
115. Independent PLU initiating cross-domain session with independent SLU	474
116. Dependent PLU initiating cross-domain session with independent SLU	474

117.	PLU initiating request for single gateway VTAM and single gateway NCP	475
118.	Independent PLU requesting session with independent SLU through single gateway VTAM and single gateway NCP	476
119.	PLU-initiated request setup queued for single gateway NCP and single gateway VTAM	477
120.	Secondary logical unit initiate (LOGON)	478
121.	Secondary logical unit initiate (REQSESS)	479
122.	Secondary logical unit initiate (INIT SELF).	480
123.	Sending an unformatted request to the SSCP	480
124.	Dependent SLU initiating a cross-domain session with application LU	481
125.	SLU initiating request for single gateway VTAM and single gateway NCP	482
126.	SLU initiating request for single gateway connecting three or more networks	483
127.	SLU initiating request for predesignated control of gateway NCP by middle host (part 1 of 2)	484
128.	SLU initiating request for predesignated control of gateway NCP by middle host (part 2 of 2)	485
129.	Third party initiating CLSDST PASS	486
130.	Third party initiating CLSDST PASS with NOTIFY	487
131.	Third party initiating request spanning three networks.	488
132.	Initiating session using VARY NET,LOGON or LOGAPPL	489
133.	Notification of PLU availability for autologon.	490
134.	Failure (CDINIT rejection) of session initiated by an SLU for single gateway VTAM and single gateway NCP	491
135.	Failure (SETCV failure) of session initiation by an SLU for single gateway VTAM and single gateway NCP	491
136.	Failure (CINIT rejection) of setup procedure initiated by an SLU for single gateway VTAM and single gateway NCP	492
137.	Deactivating a logical unit: Immediate	494
138.	Deactivating a logical unit: Forced	494
139.	Deactivating a logical unit with giveback	495
140.	Independent primary logical unit (PLU) sends BFCLEANUP for cross-domain LU-LU session with independent secondary logical unit (SLU)	495
141.	Independent primary logical unit (PLU) sends UNBIND for cross-domain LU-LU session with independent secondary logical unit (SLU)	496
142.	Primary logical unit (PLU) sends UNBIND for multiple gateway VTAMs and single gateway NCP	496
143.	Primary logical unit (PLU) sends UNBIND for single gateway VTAM and single gateway NCP	497
144.	Secondary logical unit (SLU) requests TERMINATE SELF for multiple gateway VTAMs and back-to-back gateway NCPs	498
145.	Secondary logical unit (SLU) requests TERMINATE SELF (CLEANUP) for single gateway VTAM and single gateway NCP	499
146.	Secondary logical unit (SLU) requests TERMINATE SELF for single gateway VTAM and single gateway NCP	500
147.	Active session termination of type 2.1 nodes	501
148.	Deactivating a PU acting as an adjacent link station for independent LU sessions.	501
149.	Deactivating sessions or LUs using VARY NET,TERM unconditional	502
150.	Deactivating sessions or LUs using VARY NET,TERM cleanup	503
151.	Terminating a queued session	503
152.	CLOSE ACB processing	504
153.	Deactivating an application program.	505
154.	Deactivating a CDRM: Normal.	506
155.	Deactivating a CDRM: Immediate.	507
156.	Deactivating a CDRM: Forced	507
157.	Deactivating a CDRM without affecting active sessions: Immediate	508
158.	Deactivating a CDRM without affecting active sessions: Forced	508
159.	Deactivating a CDRM on a VTAM level before V3R4.1: Forced or immediate	509
160.	SSCP-SSCP session termination causes LU-LU sessions to be broken	510
161.	Route failure in intermediate network causes termination of LU-LU sessions	511
162.	Route failure in intermediate network causes termination of SSCP-SSCP sessions	511
163.	Error recovery processing: Soft INOP	513
164.	Error recovery processing: Hard INOP	513
165.	FORWARD and DELIVER routing	514
166.	Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 PU with one link segment	514
167.	Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 PU with two link segments	515
168.	Unsolicited LPDA-2 test on permanent link error with two link segments	516
169.	Establishment of XRF primary and backup sessions.	517

170.	XRF session switch (takeover)	518
171.	Secondary logical unit initiate with USERVAR (LOGON)	519
172.	Third-party initiate (CLSDST PASS)	520
173.	CP-CP contention-winner session activation	526
174.	CP-CP contention-loser session activation	528
175.	Host CP initiating deactivation of CP-CP session.	529
176.	Remote node initiating deactivation of CP-CP session	530
177.	Activating a leased APPN node type 2.1	531
178.	Activating an APPN host-to-host channel	532
179.	Resource registration: EN to NN to CDS	534
180.	Resource registration with error recovery	535
181.	Locate resource: EN to NN	536
182.	Locate resource: EN to NN to EN	536
183.	Locate resource: EN to NN to two ENs	537
184.	Locate resource: EN to NN to NN to NN	538
185.	Locate resource: CP network broadcast initiation.	539
186.	Locate resource: EN to NN to subarea network	540
187.	Locate resource: Complex APPN network	541
188.	Locate resource: Complex APPN network using more than one CDS (part 1 of 2)	543
189.	Locate resource: Complex APPN network using more than one CDS (part 2 of 2)	544
190.	Locate resource: Network node server, NN1, of the originating logical unit (OLU) is at pre-V4R2 level (part 1 of 2)	546
191.	Locate resource: Network node server, NN1, of the originating logical unit (OLU) is at pre-V4R2 level (part 2 of 2)	547
192.	Locate resource: APPN and subarea network (part 1 of 3).	548
193.	Locate resource: APPN and subarea network (part 2 of 3).	549
194.	Locate resource: APPN and subarea network (part 3 of 3).	550
195.	Locate resource: Directory search verification reduction	551
196.	Locate resource: SLU-initiated session	552
197.	Locate resource: CP-CP session terminates	554
198.	Locate resource: Network node receives network broadcast request	555
199.	EN (PLU)—NNS...APPN network, PLU-initiated, with no queueing	558
200.	EN (PLU)—NNS...APPN network, PLU-initiated, queued by the PLU.	559
201.	EN (PLU)—NNS...APPN network, PLU-initiated, queued by the SLU	559
202.	EN (SLU)—NNS...APPN network, SLU-initiated, with no queueing	560
203.	EN (SLU)—NNS...APPN network, SLU-initiated, queued by the PLU	560
204.	APPN network...NNS—EN (SLU), PLU-initiated, no queueing	561
205.	APPN network...NNS—EN (SLU), PLU-initiated, queued by the SLU	561
206.	APPN network...NNS—EN (PLU), SLU-initiated, no queueing	562
207.	APPN network...NNS—EN (PLU), SLU-initiated, queued by the PLU.	562
208.	SA(PLU)==ICN...APPN network (SLU), DSRLIST transforming into PLU-initiated, search-only.	563
209.	SA (PLU)==ICN...APPN network (SLU), PLU-initiated, with no queueing	564
210.	SA (PLU)==ICN...APPN network (SLU), PLU-initiated, USERVAR resolution required	565
211.	SA (PLU)==ICN...APPN network (SLU), PLU-initiated, queued by the SLU (part 1 of 2)	565
212.	SA (PLU)==ICN...APPN network (SLU), PLU-initiated, queued by the SLU (part 2 of 2)	566
213.	SA (SLU)==ICN...APPN network (PLU), SLU-initiated, no queueing	567
214.	SA (SLU)==ICN...APPN network (PLU), SLU-initiated, queued by the PLU	568
215.	SA (SLU)==ICN...APPN network (PLU), autologon, PLU not available initially	569
216.	SA(PLU)==ICN...APPN network(SLU), orderly termination of active session	570
217.	APPN network(PLU)...ICN==(SA)SLU, orderly termination of active session	571
218.	SA(SLU)==ICN...APPN network(PLU), forced termination of pending active session.	571
219.	SA(PLU)==ICN...APPN network(SLU), forced termination of pending active session (PLU accessible without going into APPN)	572
220.	SA(PLU)==ICN...APPN network(SLU), forced termination of queued session	573
221.	SA(PLU)==ICN...APPN network(SLU), forced termination of queued session (PLU accessible without going into APPN)	573
222.	SA (PLU)==ICN...APPN network (SLU), session release request.	574
223.	SA (SLU)==ICN...APPN network (PLU), session release request.	574
224.	CLSDST PASS through APPN. The SLU is single-session capable.	575
225.	CLSDST PASS from APPN to subarea. The SLU is single-session capable.	576
226.	EN-NN-EN, PLU-initiated, no queueing (Including BIND flows for intermediate network node).	577

227.	Intermediate network node (INN) BIND	578
228.	APPN network (PLU)...ICN==SA(SLU), PLU-initiated, search-only flow transformed into a DSRLST	578
229.	APPN network (PLU)...ICN==SA(SLU), PLU-initiated, no queueing	579
230.	APPN network (PLU)...ICN==SA(SLU), PLU-initiated, directed search without required precomputed RSCV	580
231.	APPN network (PLU)...ICN==SA(SLU), PLU-initiated, USERVAR resolution required	581
232.	APPN network (PLU)...ICN==SA(SLU), PLU-initiated, queued by the SLU	582
233.	APPN network (SLU)...ICN==SA(PLU), SLU-initiated, no queueing	583
234.	APPN network (SLU)...ICN==SA(PLU), SLU-initiated, queued by the PLU (part 1 of 2).	584
235.	APPN network (SLU)...ICN==SA(PLU), SLU-initiated, queued by the PLU (part 2 of 2).	585
236.	APPN network (SLU)...ICN==SA(PLU), autologon (PLU not available initially)	586
237.	APPN network (PLU)...ICN==VR-based TG==ICN...APPN network (SLU), PLU-initiated	587
238.	DLUR-initiated CPSVRMGR pipe activation	589
239.	DLUS-initiated CPSVRMGR pipe activation	590
240.	Dynamic PU activation	591
241.	Dynamic registration and activation of dependent LUs.	592
242.	Activation of predefined dependent LUs	592
243.	SSCP-PU session activation race	593
244.	CPSVRMGR pipe deactivation	594
245.	Downstream PU outage	595
246.	Receipt of REQDISCONT (normal) from downstream PU	596
247.	Receipt of REQDISCONT (immediate) from downstream PU.	597
248.	Normal SSCP-PU/SSCP-LU session deactivation.	598
249.	Forced SSCP-PU/SSCP-LU session deactivation	599
250.	Giveback SSCP-PU/SSCP-LU session deactivation (ANS=STOP)	600
251.	Giveback SSCP-PU/SSCP-LU session deactivation (ANS=CONT)	601
252.	APPN PLU-initiated LU-LU session to a dependent SLU	602
253.	USS SLU-initiated LU-LU session to APPN PLU	603
254.	USS SLU-initiated LU-LU session to subarea PLU	604
255.	USS flows for LU-LU session termination	605
256.	PLU-initiated search with DLUS and DLUR within different subnetworks, PLU through the subarea (part 1 of 2)	606
257.	PLU-initiated search with DLUS and DLUR within different subnetworks, PLU through the subarea (part 2 of 2)	607
258.	PLU-initiated session with DLUS and PLU in same subnetwork and DLUR in another (part 1 of 2)	609
259.	PLU-initiated session with DLUS and PLU in same subnetwork and DLUR in another (part 2 of 2)	609
260.	SLU-initiated session with DLUS and DLUR within different subnetworks (part 1 of 2)	611
261.	SLU-initiated session with DLUS and DLUR within different subnetworks (part 2 of 2).	612
262.	An example of two Rapid-Transport Protocol (RTP) nodes with a T2.1 connection	615
263.	Two Rapid-Transport Protocol (RTP) nodes with virtual-route-based transmission group	616
264.	Rapid-Transport Protocol (RTP) connection over portion of session path	617
265.	Rapid-Transport Protocol (RTP) across composite nodes with T2.1 connection through NCP	618
266.	Rapid-Transport Protocol (RTP) across composite nodes with T2.1 connection through VTAM	619
267.	Rapid-Transport Protocol (RTP) across composite nodes with a virtual-route-based transmission group, NCP does ANR routing	620
268.	Rapid-Transport Protocol (RTP) across composite nodes with a virtual-route-based transmission group, VTAM does ANR routing	621
269.	Sample IPCS panel BLSPPRIM customization	642
270.	Addition of option 7 to the IPCS primary option menu	643
271.	Main menu for selecting dump options	643
272.	Sample ISPF panel ISR@PRIM customization	645
273.	Addition of option V to the ISPF/PDF primary option menu	646
274.	Main menu for selecting trace parameters	646

Tables

1.	Index of common problems in subarea networks	4
2.	Index of common problems in APPN networks	23
3.	Index of common problems in HPR networks	38
4.	Index of problem symptoms and associated VTAM problem types	42
5.	VIT analysis tool problems: Common symptoms and actions	45
6.	Recommended documentation for VTAM problems	50
7.	VTAM locks	74
8.	IST messages associated with CSA storage problems	97
9.	IST messages associated with private storage problems	99
10.	ABEND0AB information in LOGREC	108
11.	ABEND0AB information in a dump of SDWA	109
12.	Dumping the NCP.	186
13.	Processing externally recorded trace data	323
14.	Printing external trace entries	323
15.	Symbols and numbers for formatting and printing VTAM traces	324
16.	Fields in VTAM trace output	325
17.	Fields in the buffer contents trace	334
18.	Fields in the SMS trace	350
19.	Location of TPUT (outbound) error	351
20.	Location of TGET (inbound) error.	351
21.	Fields in the transmission group trace	361
22.	Boolean expression operators in order of precedence	391
23.	VTAM channel commands for type 2 and type 4 physical units.	414
24.	PUNS-related channel programs	419
25.	Channel commands for channel-to-channel adapters	422
26.	Channel command words for channel-attached non-SNA 3270 devices	441
27.	Write data channel program	442
28.	Read Modified channel program	443
29.	Read buffer channel program	443
30.	Erase/Write channel program	444
31.	Erase/Write Alternate channel program.	445
32.	Erase All Unprotected channel program.	445
33.	Index of network flows	447
34.	Index of generic BIND (GBIND) AMRU flows	451
35.	Index of resource activation flows.	453
36.	Index of session establishment flows.	466
37.	Index of deactivation and session termination flows	492
38.	Index of error detection and recovery and SSCP management services flows	512
39.	Index of APPN flows	521
40.	Index of CP-CP session flows	525
41.	Index of directory services flows	533
42.	Index of LU-LU session flows	556
43.	Index of dependent LU server flows	587
44.	Index of High-Performance Routing flows	614
45.	Control point/control block operation codes (CPCBOPC)	623
46.	Control block ID codes	637
47.	Target data sets for dump and trace tools	641
48.	Related information on problem topics in other libraries	649

About this document

This document is intended to help system programmers in a VTAM[®] environment to diagnose problems with the VTAM program. Use the document to isolate and identify problems with your VTAM network and to collect appropriate documentation to resolve network problems.

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.

Who should read this document

System programmers should use this document to analyze a VTAM problem, classify the problem as a specific type, and provide information about the problem to an IBM[®] Support Center representative.

How this document is organized

This document is organized into the following parts:

- Part 1, "Diagnostic techniques," on page 1 describes how to identify a problem.
- Part 2, "Diagnostic procedures," on page 121 describes how to use diagnostic procedures.
- Appendix A, "Channel programs," on page 413, Appendix B, "Network flows," on page 447, Appendix B, "Network flows," on page 447, Appendix D, "Control point/control block (CPCB) operation codes," on page 623, Appendix E, "Storage and control block ID codes," on page 637, Appendix F, "Installing dump analysis and VIT analysis tools," on page 641, Appendix G, "Problem topics in other libraries," on page 649, Appendix H, "Architectural specifications," on page 653, and Appendix I, "Accessibility," on page 655 provide additional information for this document.

How to use this document

You should be familiar with the service aids for VTAM and the procedures for reporting problems to an IBM Support Center representative.

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 683.

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 term RDMA network interface card (RNIC) that is used to refer to the IBM 10GbE RoCE Express feature.

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 section describes how to read the syntax diagrams used in this book.

- Read the diagrams from left-to-right, top-to-bottom, following the main path line. Each diagram begins on the left with double arrowheads (▶▶) and ends on the right with two arrowheads facing each other (◀◀).



- 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 and values appear on the main path line.




You must code required operands and values.

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



- Optional operands and values appear below the main path line.



You can choose not to code optional operands and values.


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



- An arrow returning to the left above an operand or value on the main path line means that the operand or value can be repeated. The comma means that each operand or value must be separated from the next by a comma.



- 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.



- A word in all uppercase is an operand or value you must spell exactly as shown. In this example, you must code *OPERAND*.

Note: VTAM and IP commands are not case sensitive. You can code them in uppercase or lowercase. If the operand is shown in both uppercase and lowercase, the uppercase portion is the abbreviation (for example, OPERand).



If an operand or value can be abbreviated, the abbreviation is described in the text associated with the syntax diagram.

- 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)*.



- 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 and values appear above the main path line. VTAM uses the default if you omit the operand entirely.



- 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.



- References to syntax notes appear as numbers enclosed in parentheses above the line. Do not code the parentheses or the number.



Notes:

- 1 An example of a syntax note.
- 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 669, 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
<i>IBM System z® Redbooks Collection</i>	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 zSeries® platform professional are identified by their authors and are included in this collection. The zSeries 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, including the z/OS Information Center, see www.ibm.com/systems/z/os/zos/bkserv/.

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

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

Title	Number
<i>DNS and BIND</i> , Fifth Edition, O'Reilly Media, 2006	ISBN 13: 978-0596100575
<i>Routing in the Internet</i> , Second Edition, Christian Huitema (Prentice Hall 1999)	ISBN 13: 978-0130226471
<i>sendmail</i> , Fourth Edition, Bryan Costales, Claus Assmann, George Jansen, and Gregory Shapiro, O'Reilly Media, 2007	ISBN 13: 978-0596510299
<i>SNA Formats</i>	GA27-3136
<i>TCP/IP Illustrated, Volume 1: The Protocols</i> , W. Richard Stevens, Addison-Wesley Professional, 1994	ISBN 13: 978-0201633467
<i>TCP/IP Illustrated, Volume 2: The Implementation</i> , Gary R. Wright and W. Richard Stevens, Addison-Wesley Professional, 1995	ISBN 13: 978-0201633542
<i>TCP/IP Illustrated, Volume 3: TCP for Transactions, HTTP, NNTP, and the UNIX Domain Protocols</i> , W. Richard Stevens, Addison-Wesley Professional, 1996	ISBN 13: 978-0201634952
<i>TCP/IP Tutorial and Technical Overview</i>	GG24-3376
<i>Understanding LDAP</i>	SG24-4986
z/OS Cryptographic Services System SSL Programming	SC24-5901
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 196, System z10, System z9 and eServer zSeries 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
<i>IBM z/OS V1R13 Communications Server TCP/IP Implementation, Volume 1: Base Functions, Connectivity, and Routing</i>	SG24-7996
<i>IBM z/OS V1R13 Communications Server TCP/IP Implementation, Volume 2: Standard Applications</i>	SG24-7997
<i>IBM z/OS V1R13 Communications Server TCP/IP Implementation, Volume 3: High Availability, Scalability, and Performance</i>	SG24-7998
<i>IBM z/OS V1R13 Communications Server TCP/IP Implementation, Volume 4: Security and Policy-Based Networking</i>	SG24-7999
<i>IBM Communication Controller Migration Guide</i>	SG24-6298
<i>IP Network Design Guide</i>	SG24-2580
<i>Managing OS/390® TCP/IP with SNMP</i>	SG24-5866
<i>Migrating Subarea Networks to an IP Infrastructure Using Enterprise Extender</i>	SG24-5957
<i>SecureWay Communications Server for OS/390 V2R8 TCP/IP: Guide to Enhancements</i>	SG24-5631
<i>SNA and TCP/IP Integration</i>	SG24-5291
<i>TCP/IP in a Sysplex</i>	SG24-5235
<i>TCP/IP Tutorial and Technical Overview</i>	GG24-3376
<i>Threadsafe Considerations for CICS</i>	SG24-6351

Where to find related information on the Internet

z/OS

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

<http://www.ibm.com/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/atmastr.nsf>

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 III1334.

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://publib.boulder.ibm.com/infocenter/zos/basics/index.jsp>

Summary of changes

This section describes the release enhancements that were made.

New in 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. Diagnostic techniques

Chapter 1. Diagnosing VTAM problems: Where to begin

This information includes the following topics:

- To help you determine the source of your problem, see “Determining whether the problem is VTAM or non-VTAM.”
- To compare your problem to a list of common problems that have been identified by the IBM Support Center, see “Common problems in subarea networks” on page 4, “Common problems in APPN networks” on page 22, and “Common problems in HPR networks” on page 38.

For additional information, see “Common symptoms and associated VTAM problem types” on page 42.

- If you are having problems with the trace or dump analysis tools, see “VTAM internal trace (VIT) analysis tool problems” on page 44 and “VTAM dump analysis tool problems” on page 47.
- To gather additional documentation to help you to solve your problem, see “Recommended documentation for VTAM problems” on page 48.
- To prepare your documentation for submission to the IBM Support Center, see “Submitting documentation on tape” on page 54.

Determining whether the problem is VTAM or non-VTAM

Problems can be classified into two types:

VTAM problems

These are problems that occur in the VTAM program.

Non-VTAM problems

These are problems that occur because of changes to your network or problems caused by other application programs or software in the network, such as a network control program (NCP) or a local area network (LAN).

If you did any of the following actions, the problem might be in your network setup, in your configuration, or in another IBM product:

- Did you modify an application program that has run without problems in the past?
- Did you modify a product exit routine that has run without problems in the past?
- Did you change the processing environment? For example, did you introduce a new host processor or communication controller?
- Did you modify the operating system, or did you install a new release of the operating system?
- Did you add a new terminal to your VTAM network that had incorrect features or incorrect Request for Engineering Activity (REA) and Engineering Change (EC) levels?
- Did you attach a link?
- Did you set switches at a terminal?
- Did you initialize link parameters for a programmable controller?

- Did you modify installation-provided VTAM tables? For example, did you modify logmode, Class of Service (CoS), or unformatted system services (USS) tables?
- Did you modify VTAM messages?

If you cannot resolve the problem on your own:

- Compare your problem to the examples in “Common problems in subarea networks” and “Common problems in APPN networks” on page 22.
- Check Table 4 on page 42 for your problem symptom.
- Follow the instructions in “Recommended documentation for VTAM problems” on page 48. To gather additional information, see Chapter 2, “Collecting documentation for specific types of problems,” on page 57.
- For non-VTAM problems, call your IBM branch office. For suspected VTAM problems, do either of the following steps:
 - If you have access to IBMLink, search for known problems in this area. If no applicable matches are found, report the problem to IBM by using the electronic technical report (ETR) option on IBMLink.
 - Contact the IBM Software Support Center at 1-800-IBM-SERV.

If a reported problem is a hardware, network definition, or user definition error, the IBM Support Center representative creates an ASKQ item for VTAM. The ASKQ item includes the solution for the problem and can be found in the problem determination database (PDDDB).

Common problems and symptoms

This topic contains the following information:

- “Common problems in subarea networks”
- “Common problems in APPN networks” on page 22
- “Common problems in HPR networks” on page 38
- “Common symptoms and associated VTAM problem types” on page 42

Even in cases when a VTAM problem has the same symptoms as a non-VTAM problem, by studying similar examples in this topic, you may be better prepared when you contact the IBM Support Center.

If you have access to a software support database, you can search for your problem in that database and apply any recommended correction.

Common problems in subarea networks

Table 1 includes a brief description of several common problems that occur in subarea networks. For additional information, see the page indicated.

Table 1. Index of common problems in subarea networks

Problem	See page
Abend 0C4 in ISTTSCPF when move character instruction processed (reason code 4, 10, or 11)	“Abend 0C4 occurs in ISTTSCPF” on page 8
Abend in user exit ISTECCS, ISTECCSD, ISTECCAA, or ISTECCVR with a dump taken by VTAM ESTAE ISTIECXT, ISTECEX, or ISTCSCSD	“Abend in user exit ISTECCS, ISTECCSD, ISTECCAA, or ISTECCVR” on page 9

Table 1. Index of common problems in subarea networks (continued)

Problem	See page
Activating an NCP and resources fail	"Sense codes 10030000 and 08090000 received when activating an NCP" on page 22
APPL-APPL storage expansion failure with messages IST154I, IST561I, IST999E, IST566I, and IST930I	"Messages IST154I, IST561I, IST999E, IST566I, and IST930I received for APPL-APPL storage expansion failure" on page 18
BIND failure with message IST663I (sense code 083500xx or 08210000) and USSMSG07	"Sense code 083500xx or 08210000 received with BIND failure" on page 20
CICS [®] logon problem with message IST663I (sense code 08210000)	"Message IST663I (sense code 08210000) and cannot log on to CICS" on page 17
CINIT failure with message IST663I (sense code 08010000)	"Sense code 08010000 received with CINIT failure" on page 20
DSRLST pending condition and CD DSEARCH PENDING in message IST530I or message IST1278I	"Message IST530I or IST1278I received with pending DSRLST condition" on page 16
LU hung in PNFYx state.	"LU hung in PNFYx state" on page 9
Message IKT029I with return code 061001 (TSO logon failure from a session manager application program)	"Message IKT029I (return code 061001) received with TSO logon failure" on page 10
Message IST259I and sessions end unexpectedly for a terminal, PU, line, or NCP.	"Message IST259I received and sessions end unexpectedly" on page 11
Message IST259I indicating that an INOP RU was received for a link problem	"Message IST259I received with INOP RU" on page 10
Message IST264I (required CoS entry undefined) and message IST663I with sense code 08610000	"Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry" on page 17
Message IST467I with error type 05, 07, 08, or 0B during activation of a PU	"Message IST467I received with CONTACTED ERROR TYPE 05, 07, 08, or 0B" on page 12
Message IST530I (DSRLST pending condition)	"Message IST530I or IST1278I received with pending DSRLST condition" on page 16
Message IST530I (GUNBIND pending)	"Message IST530I or IST1278I received with GUNBIND PENDING or session hangs in PSESEND state" on page 15

Table 1. Index of common problems in subarea networks (continued)

Problem	See page
Message IST530I (NMVT PENDING)	"Message IST530I or IST1278I received with NMVT PENDING" on page 16
Message IST663I with sense code 08010000 (CINIT failure)	"Sense code 08010000 received with CINIT failure" on page 20
Message IST663I with sense code 08210000 (CICS logon problem)	"Message IST663I (sense code 08210000) and cannot log on to CICS" on page 17
Message IST663I (sense code 083500xx or 08210000) and USSMSG07 (BIND failure)	"Sense code 083500xx or 08210000 received with BIND failure" on page 20
Message IST663I (sense code 08610000) and message IST264I (required CoS entry undefined)	"Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry" on page 17
Message IST1278I (DSRLST pending condition)	"Message IST530I or IST1278I received with pending DSRLST condition" on page 16
Message IST1278I (GUNBIND pending)	"Message IST530I or IST1278I received with GUNBIND PENDING or session hangs in PSESEND state" on page 15
Message IST1278I (NMVT PENDING)	"Message IST530I or IST1278I received with NMVT PENDING" on page 16
Messages IST154I, IST561I, IST999E, IST566I, and IST930I (APPL-APPL storage expansion failure)	"Messages IST154I, IST561I, IST999E, IST566I, and IST930I received for APPL-APPL storage expansion failure" on page 18
PNFYx resource state and LU hung	"LU hung in PNFYx state" on page 9
PSESEND session termination state and hung session	"Message IST530I or IST1278I received with GUNBIND PENDING or session hangs in PSESEND state" on page 15
Resources fail when activating an NCP (sense codes 10030000 and 08090000).	"Sense codes 10030000 and 08090000 received when activating an NCP" on page 22
Sense code 08010000 with message IST663I (CINIT failure)	"Sense code 08010000 received with CINIT failure" on page 20

Table 1. Index of common problems in subarea networks (continued)

Problem	See page
Sense code 08210000 with message IST663I (CICS logon problem)	"Message IST663I (sense code 08210000) and cannot log on to CICS" on page 17
Sense code 083500xx or 08210000 (message IST663I) and USSMSG07 (BIND failure)	"Sense code 083500xx or 08210000 received with BIND failure" on page 20
Sense code 08610000 (message IST663I) and message IST264I (required CoS entry undefined)	"Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry" on page 17
Sense code 08610000 (message IST663I) and message IST264I (required CoS entry undefined)	"Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry" on page 17
Sense code 0888000x with session failure	"Session failure with sense code 0888000x" on page 22
Sense code 800A0000 or no message with session failure	"Sense code 800A0000 or no message, and sessions end unexpectedly" on page 21
Sense code 80130104 and path problems	"Example: Solving path problems" on page 157
Sense codes 10030000 and 08090000. (Resources fail when activating an NCP.)	"Sense codes 10030000 and 08090000 received when activating an NCP" on page 22
Session fails with sense code 0888000x.	"Session failure with sense code 0888000x" on page 22
Session hung with PSESEND session termination state.	"Message IST530I or IST1278I received with GUNBIND PENDING or session hangs in PSESEND state" on page 15
Sessions end unexpectedly for a terminal, PU, line, or NCP with message IST259I.	"Message IST259I received and sessions end unexpectedly" on page 11
Sessions end with no message or sense code 800A0000.	"Sense code 800A0000 or no message, and sessions end unexpectedly" on page 21
Storage problem	"Storage problem procedure" on page 97
TSO application program receives partial input for the TGET macroinstruction.	"Partial input for TGET received by TSO" on page 19
TSO logon failure from a session manager application program with message IKT029I and return code 061001	"Message IKT029I (return code 061001) received with TSO logon failure" on page 10

Table 1. Index of common problems in subarea networks (continued)

Problem	See page
USS message USSMSG07 and message IST663I with sense code 083500xx or 08210000 (BIND failure)	"Sense code 083500xx or 08210000 received with BIND failure" on page 20
VTAM trace records were expected but are not in the GTF trace data set.	"Missing VTAM trace records" on page 19

Descriptions of common problems in subarea networks

This topic includes examples of common problems in subarea networks. See Table 1 on page 4 for an index of these problems.

Abend 0C4 occurs in ISTTSCPF

Problem statement

An abend 0C4 with a reason code of 4, 10, or 11 occurs in all levels of the module ISTTSCPF when a move character instruction is processed.

Common symptoms

The USS message contains unexpected or extraneous characters, but not all USS messages are affected. Devices might not activate correctly or might fail to log on correctly. An abend 0Cx might occur in module ISTTSCPF.

Probable cause

The USSTAB is incorrectly defined. In this module, register 4 points to storage that may not be paged-in (reason code 10 or 11) or to storage that should not be accessed (reason code 4).

The TSCB contains a data length field (TSCDATALN) equal to the length of the USSMSG text plus the length field itself. Subtracting the value in register 4 from the starting address of the USSMSG table shows that only the value of TSCDATALN (minus 2 bytes) was moved because the full amount of storage that was referenced was not paged-in.

User response

Code only the length of the USSMSG entry in the USSMSG table. Do not include the size of the length field.

•

Problem statement

An abend 0C4 with a reason code of 4, 10, or 11 occurs in all levels of the module ISTTSCPF when a move character instruction is processed.

Common symptoms

The USS message contains unexpected or extraneous characters, but not all USS messages are affected. Devices might not activate correctly or might fail to log on correctly. An abend 0Cx might occur in module ISTTSCPF.

Probable cause

The USSTAB is incorrectly defined. In this module, register 4 points to storage that may not be paged-in (reason code 10 or 11) or to storage that should not be accessed (reason code 4).

The TSCB contains a data length field (TSCDATALN) equal to the length of the USSMSG text plus the length field itself. Subtracting the value in register 4 from the starting address of the USSMSG table shows that

only the value of TSCDATLN (minus 2 bytes) was moved because the full amount of storage that was referenced was not paged-in.

User response

Code only the length of the USSMSG entry in the USSMSG table. Do not include the size of the length field.

Abend in user exit ISTECCS, ISTECCSD, ISTECCAA, or ISTECCVR

Problem statement

A failure occurred when running VTAM Exit Facility Subtask. If an ABEND code occurs in user exit ISTECCS, ISTECCSD, ISTECCAA, or ISTECCVR, a VTAM ESTAE will attempt to take a dump for the abnormally ending user exit.

Common symptoms

An ABEND code occurred. If VTAM attempts to issue the SDUMPX, message IST413I will be issued. If SDUMPX is issued but the dump fails, message IST257I will give the reason code for the dump failure.

Note that the VRDATA KEY=DAE macro is issued before issuing the SDUMPX to make the dump eligible for MVS DAE (Dump Analysis and Elimination) dump suppression. Thus, if the user has DAE in effect at the time of the errors, duplicate dumps with matching symptoms will be suppressed by the MVS DAE facility, and VTAM will issue the following message:

```
IST257I VTAM SDUMP FAILED WITH RETURN CODE 08 REASON X'0B'
```

Probable cause

The following list shows some of the common problems with user-written exits:

- The pointer to the VTAM Exit Services parameter list was not valid when the session management exit called VTAM EXIT Services.
- The pointer to the EXMPL (the pointer to the input parameter list in the VTAM Exit Services parameter list) was nonzero, but was not valid when the exit called VTAM Exit Services.
- The pointer to the message text in the EXMPL was nonzero, but was not valid when the exit called VTAM Exit Services.
- Some portion of the message text could not be accessed by VTAM Exit Services. (For example, the session management exit passed a message length in the EXMPL that exceeded the storage area owned by the Session Management Exit.)

User response

Consult the symptom string or PSW and registers, or both, at the time of ABEND. Consult any related storage or addresses in the dump.

LU hung in PNFYx state

Problem statement

An LU can hang in a PNFYx state if the application program does not issue the CLSDST macroinstruction when a LOSTERM user exit routine is scheduled.

Common symptoms

An LU is hung in a PNFYx state. The LU is unable to log on to an application.

Probable cause

An application program failed to issue the CLSDST macroinstruction when the LOSTERM user exit routine was scheduled with a reason code indicating that the CLSDST macroinstruction should be issued.

User response

Check with the owner of the application program for known problems. A VTAM internal trace with the application interface (API) option active indicates whether the LOSTERM user exit was scheduled and the reason code that was passed. The API trace option can also be used to determine if the CLSDST was issued by the application in response to a LOSTERM user exit.

Notes:

1. See "PNFYx status" in "Using the VARY INACT,FORCE command" on page 80 for information on the VARY INACT,FORCE command and PNFYx status.
2. Coding the LOSTERM parameter on the APPL definition statement allows you to recover this type of hung resource without having to cancel the application. See the z/OS Communications Server: SNA Programming for more information.
3. The information in this problem description is from information APAR II00757. See that APAR for additional information.

Message IKT029I (return code 061001) received with TSO logon failure**Problem statement**

Cannot log on to TSO from a session manager application.

Common symptoms

The following message is displayed:

```
IKT029I RC= 061001 SENSE= code TERMINAL termid ABOUT TO BE RELEASED BY VTAM
```

Probable cause

Either the application or D/T8100 expects the first BIND from TSO to be from terminal control address space (TCAS). TCAS will send only the BIND for the TSO subapplication program (TSOxxxx).

User response

If the secondary logical unit (SLU) does not support this type of session initiation, specify FASTPASS=NO on the SLU definition statement to force TCAS to send a BIND to the SLU before the TSO subapplication program sends its own BIND.

Message IST259I received with INOP RU**Problem statement**

Message IST259I is generated by the inoperative (INOP) RU processor. The INOP RU is generated by the data link control (DLC) component for the subarea controlling the link, either intermediate network node (INN) or route extension (REX).

Common symptoms

The following message is displayed:

```
IST259I INOP RECEIVED FOR nodename CODE = code
```

If the link is INN, an ER.INOP will also flow, producing a series of explicit route (ER) or virtual route (VR) failure messages. This leads to an incorrect diagnosis when you do not associate the ER.INOP with the link or link station INOP.

If message IST259I contains the name of a channel-attached NCP or a local device, message IOS000I might accompany the failure.

If message IST259I contains the name of a channel-attached 3172 device, messages IST1411I, IST1412I, or IST1430I will indicate the reason for the INOP. For more information, see *z/OS Communications Server: SNA Messages*.

Some local SNA controllers require I/O buffer size to be an even number. For example, if an odd number is coded for a 3174, message IST259I with CODE=01 will be displayed at activation.

Probable Cause

NCP link

This is a communication facility problem. Either the retry limit is exhausted, a negative acknowledgment is received for an SDLC transmission, a modem error occurred, or a link failure occurred.

Channel link

Either a data transfer count mismatch occurred, an NCP abend has occurred, or the NCP was reloaded by another host.

User Response

Trace the link.

List the system LOGREC to obtain the data from the record management statistics (RECMS) that accompany an INOP originating in an NCP node. The RECMS identifies the error that produced the INOP. Use environmental record editing and printing (EREP) to print the LOGREC records. Use the network problem determination application (NPDA) to interpret the RECMS record.

VTAM does not generate the RECMS for channel link and link station failures. The LOGREC entry for a local device will contain only statistical data.

Correct the error condition.

Note: The IBM Support Center representative can only suggest that you list LOGREC and assist you with interpreting the record.

If a channel-attached SNA device (NCP or cluster controller) is experiencing the INOPs at a regular or predictable interval (for example, every hour), the problem could be that the VTAM ERP routine has been deleted. Verify that CSECT ISTZBM0K in load module IGE0004 (LPALIB) has not been deleted.

Message IST259I received and sessions end unexpectedly

Problem statement

One or more sessions have ended unexpectedly, and a terminal, PU, line, or NCP is in a wait state.

Common symptoms

The following message is displayed:

```
IST259I  INOP RECEIVED FOR nodename CODE = code [text]
```


Probable cause

- If the node is an NCP, the NCP detected an error and generated the INOP RU message.
- If the node is a channel-to-channel (CTC) link or a CTC link station, VTAM detected an error from an IO operation and generated the INOP RU message.
- If the node is a local attachment device, VTAM detected an IO error and generated the INOP RU message.

User response

- For an explanation of the code in IST259I, see the description of the message in z/OS Communications Server: SNA Messages.
- Check the system log for system (IOS) error messages that contain status information.
- If the NetView[®] program is installed, check NPDA for logged errors.
- Run EREP against LOGREC, and check for errors related to the device.

Note: This information should identify the component causing the error. Contact the appropriate service organization for help with a specific component problem.

Message IST467I received with CONTACTED ERROR TYPE 05, 07, 08, or 0B

Problem statement

Message IST467I is received with contacted error type 05, 07, 08, or 0B during activation of a resource. The message indicates that the XID was rejected by the PU.

Common symptoms

IST467I is the first in a group of messages. The exchange ID (XID) received by VTAM is shown in messages IST1574I and IST1580I. The XID sent by VTAM is shown in messages IST1574I and IST1586I. Compare the XIDs to determine why the PU rejected the XID.

Sample XIDs from an IST467I message group:

```
*****
* The following is for XID format 2. All *
* references to bytes and bits are in hex. *
*****
XID1 (Received from the NCP) =
  242AFF0 00000000 00080000 00010000
  00035007 D5C3D7D3 D6C1C440 80000203
  002A05F3 00800000 0000
XID2 (Sent to the NCP by VTAM) =
  242AFF0 00000000 200800F9 DE010000
  00010000 40404040 40404040 81000200
  002A05F3 00000000 0000
```

Note: If *type* is 0B in message IST467I, additional error information may be contained in a CV X'22' appended to the end of the XID. See the **User Response** for an example.

Probable cause

Note: The explanations that follow cover more than the single error that the sample XID1 and XID2 represent.

The first digit of the XID is the format. In the preceding sample XIDs, the format is format 2.

Byte X'12' of the XID1 received from the NCP is the error byte.

•

- Bit 0** Reserved (unused)
- Bit 1** Received XID unacceptable
- Bit 2** Incompatible
- Bit 3** Transmission group (TG) undefined

Bit 1 of byte X'12' is set for the following reasons:

1. The XID2 at displacement X'00' was not equal to 24 or 25.
2. The XID2 at displacement X'08' was not equal to 20.
3. The XID2 at displacement X'13' was not equal to 00.
4. The XID2 at displacement X'1E' was not equal to 02.
5. Depending on the release of NCP you have:
 - a. *For NCP V4R3, V5R2, and higher:* The XID2 from VTAM at displacement X'0B'—X'0C' is less than 1296 decimal (X'0510'). The value in this field is the result of MAXBFRU from the HOST macro times the IOBUF buffer size in the VTAM start list. This error is set only when the NCP definitions have specified the HOST connection using GROUP LNCTL=CA.
 - b. *For NCP before V4R3 and V5R2:* The XID2 from VTAM at displacement X'0B'—X'0C' is less than (XID plus X'20'—X'21') times (XID plus X'22'—X'23') minus (XID plus X'24') in the XID sent by NCP. This result corresponds to the value specified on the MAXBFRU operand times the value specified on the UNITSZ operand minus the value specified on the BFRPAD operand.
6. There is no path to the subarea number defined at X'11' in the XID2 that uses this connection.
7. Bit 2 of byte X'12' is set because the received XID1 at X'25' is not equal to X'20', and an existing connection exists with the origin subarea.
8. Bit 3 of byte X'12' is set because either the TG number in the XID2 at displacement X'0D' or the subarea from the XID2 at X'11' is unknown to the NCP.

User Response

Reasons 2, 3, and 4 should not occur, but should help to verify XID offsets.

For reason 5, see the *z/OS Communications Server: SNA Resource Definition Reference* and the *NCP, SSP, and EP Resource Definition Reference* regarding specification of buffer sizes.

For reason 6, see the *z/OS Communications Server: SNA Resource Definition Reference* and the *NCP, SSP, and EP Resource Definition Reference* regarding the definition of PATH statements. Also, transmission group (TG) mismatch could cause the problem. A TG mismatch could occur, for example, if an NCP is attached as a CA major node, and TG=ANY is coded in the CA major node in VTAM, and TG=ANY is coded on the NCP line definition for this attachment.

For reason 7, see the *z/OS Communications Server: SNA Resource Definition Reference* regarding the use of the CHANCON parameter of the PCCU macro.

For reason 8 on page 13, see the *NCP, SSP, and EP Resource Definition Reference* regarding the use of the CANETID parameter on the BUILD macro. Verify that it is coded correctly for each network in which it is assigned.

Note:

1. In the sample XID1 given above, byte X'12' contains the value X'50'. Bit 2 and bit 3 indicate that the XID2 was unacceptable and that the transmission group was not defined. The problem in this case was that there were no PATH definition statements defined in the NCP for the host subarea.
2. The contacted error type 05 can also be posted if the NETIDs in the 2 XIDs do not match.

The NETID will appear in a CV X'12' at the end of the XIDs. If they do not match, correct the NETID operand on the PU definition statement in one or both PU definitions.

3. A contacted error type 05 may occur if a channel-to-channel connection is defined between two VTAM systems that have the same subarea. Subarea numbers must be unique.

In this case, the 4 bytes starting at offset X'E' in XID1 and XID2 will be the same. XID offset X'E' contains the subarea numbers. XID1 contains the subarea number of the receiver, and XID2 contains the subarea number of the sender.

The following example is for type 0B when a CV X'22' is appended to the end of the XID.

```

IST467I CONTACTED ERROR TYPE 0B FOR ID = AHHCPU1
IST1580I XID RECEIVED BY VTAM:
IST1581I +000 348AFFF0 99260000 10CB4100 00000080 *...0.....
IST1581I +010 00060530 0000000E 09F1C9E2 E3D7E4E2 *.....1ISTPUS
IST1581I +020 40400E0C F4D5C5E3 C14BE2E2 C3D7F1C1 * ..4NETA.SSCP1A
IST1581I +030 0E08F7C1 C8C8C3D7 E4F14609 09800000 *..7AHHCPU1.....
IST1581I +040 00000000 01103A00 2311040E 02F5F6F9 *.....569
IST1581I +050 F5F1F1F7 F0F1F8F0 F10804F0 F4F0F4F0 *511701801..04040
IST1581I +060 F00A06C1 C3C661E5 E3C1D416 11011300 *0..ACF/VTAM.....
IST1581I +070 11F9F0F2 F1000000 0000F0F1 F3F2F0F8 *.9021.....013208
IST1581I +080 F2220700 09040000 0000 *2.....
IST1582I CONTROL VECTOR 22 ANALYSIS:
IST1583I BYTE OFFSET OF FIRST BYTE IN ERROR = X'0009'
IST1584I BIT OFFSET OF FIRST BIT IN ERROR = X'04'
IST1586I XID SENT BY VTAM:
IST1581I +000 34B1FFF0 992B0000 10F74100 00000080 *...0.....7.....
IST1581I +010 15060530 0010000E 09F1C9E2 E3D7E4E2 *.....1ISTPUS
IST1581I +020 40400E0C F4D5C5E3 C14BE2E2 C3D7F2C1 * ..4NETA.SSCP2A
IST1581I +030 0E08F7C1 C8C8C3D7 E4F14609 09801500 *..7AHHCPU1.....
IST1581I +040 00000000 02612E30 00088001 00280000 *...../.....
IST1581I +050 00002381 141D0000 007800AC DED371AC *.....L..
IST1581I +060 DED37108 D4000000 00000000 08D20000 *.L..M.....K..
IST1581I +070 00000000 00103A00 2311040E 02F5F6F9 *.....569
IST1581I +080 F5F1F1F7 F0F1F8F0 F10804F0 F4F0F4F0 *511701801..04040
IST1581I +090 F00A06C1 C3C661E5 E3C1D416 11011300 *0..ACF/VTAM.....
IST1581I +0A0 11F9F0F2 F1000000 0000F0F1 F3F2F0F8 *.9021.....013208
IST1581I +0B0 F2 *2
IST314I END

```

see the description of message IST467I in z/OS Communications Server: SNA Messages for additional information.

Message IST530I or IST1278I received with GUNBIND PENDING or session hangs in PSESEND state

Problem statement

A GUNBIND PENDING message is received at logoff time in a cross-domain environment (if IOPD is specified or defaulted in the VTAM start options), or the session hangs in PSESEND session termination state.

Common symptoms

- Message IST530I or IST1278I:
GUNBIND PENDING FROM app1name TO LU
- Message IST530I or IST1278I
GUNBIND PENDING FROM VTAM TO LU
- The session displays PSESEND as the session termination state.

Probable cause

- The application did not issue a CLSDST macroinstruction.
- The device did not respond to the UNBIND request or returned a response that was incorrect or not valid.
- A virtual route between the primary logical unit (PLU) subarea and the secondary logical unit (SLU) subarea is held or blocked.
- The network ID defined in the NCP does not match the network ID coded in the VTAM start options.

User response

- Enter (on the terminal owning the host) and note the status:
D NET, ID=devicename, E
- If the session termination state is PSESEND, enter:
D NET, SESSIONS, SID=*sid*

(where the SID is that of the PSESEND session in the IST635I message group)

This display will show which session partner is withholding the session end signal to complete the session termination.

- If a SESSEND is needed from the PLU, VTAM is waiting for a CLSDST macroinstruction to be issued.
- If a SESSEND is needed from the SLU, there is usually a problem in a network element, such as the host VTAM, NCP, or SLU.
- If the device is remote and hangs in PSESEND session termination state at logoff, start the following trace (in the device-owning host), and trace a logon and logoff:

```
F NET, TRACE, TYPE=BUF, ID=devicename
```

Check the X'15' vector in the SESSST and SESSEND RUs (if they are present) to see if the network ID matches the network ID coded in the VTAM start options.

- Enter (on the application-owning host) and note the status:
D NET, ID=app1name, E

This indicates whether other sessions are affected. If the application name has many sessions, this display output can be very large.

```
D NET, ID=devicename, E
```

This indicates whether the session status matches the session status in the device host.

D NET,ROUTE,DESTSUB=device_subarea_number,TEST=YES

This indicates that a virtual route is held or blocked.

D NET,TERM,SID=sid,TYPE=FORCE

This may help expedite session termination.

- If it is suspected that no CLSDST macroinstruction is being issued, a buffer trace of the application and a VTAM internal trace with MODE=EXT,OPT=API specified may be needed to verify:
 - That the application was notified of session termination
 - Which exit was scheduled
 - What actions or commands were issued (if any) by the application

Message IST530I or IST1278I received with NMVT PENDING

Problem statement

Message IST530I or IST1278I is issued for a PU even though the NetView program (if installed) or the System p[®] network management program for System p devices receives session awareness (SAW) data for an SNA device.

Common symptoms

Message IST530I or IST1278I is issued each time the IOPD timer expires. For additional information, see the message descriptions in z/OS Communications Server: SNA Messages.

Probable cause

The device is not real-time monitor capable. This means that the device did not process the response and return the requested information properly to the NetView program for most devices, or to Network Management/6000 for System p devices. A microcode change is needed to permanently resolve this problem.

User response

You can prevent this problem by pointing the device to a KCLASS and using a SAW data filter to stop VTAM from attempting to collect the data.

Note: See the z/OS Communications Server: SNA Network Implementation Guide and z/OS Communications Server: SNA Resource Definition Reference for details on how to code a SAW data filter.

Message IST530I or IST1278I received with pending DSRLST condition

Problem statement

A DSRLST PENDING message is received. Message IST530I or IST1278I is issued with CD DSEARCH PENDING FROM netid TO netid.

Common symptoms

Message IST530I or IST1278I is issued for the application.

Probable cause

- The ADJSSCP table was not coded; the ADJSSCP table is coded incorrectly; or the IOINT value is too low.
- The start option DYNASSCP and the ADJSSCP table are not correctly tuned.

User response

- To identify the ADJSSCP, enter (with or without a NETID operand):
D NET,ADJSSCPS

- To determine the current value of IOINT, enter:
D NET,VTAMOPTS,OPTIONS=IOINT
- To identify the ADJSSCP, enter (with or without a NETID operand):
D NET,ADJSSCPS
- To determine the DYNASSCP value, enter one of the following codes and note the DYNASSCP value specified:
D NET,VTAMOPTS,OPTIONS=*
D NET,VTAMOPTS,OPTIONS=DYNASSCP

Message IST663I (sense code 08210000) and cannot log on to CICS

Problem statement

Sessions cannot log on to CICS.

Common symptoms

The message IST663I CINIT REQUEST FROM *adjnode* FAILED, SENSE=08210000 is received.

Probable cause

When running CICS with AUTO-INSTALLATION, the terminal definition in the terminal control table terminal entry (TCTTE) must match the VTAM LOGMODE definition statement for the device.

User response

Either change the VTAM LOGMODE definition statement to match the CICS TCTTE, or code LOGMODE=0 in the TCTTE. Adding LOGMODE=0 to the TCTTE forces CICS to use VTAM's LOGMODE definition statement for this session.

Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry

Problem statement

A required CoS entry is UNDEFINED.

Common symptoms

The following messages are received:

```
IST663I request REQUEST FAILED, SENSE=08610000
IST264I REQUIRED COS luname UNDEFINED
HASP208 LOSTTERM SCHEDULED SNA, VTAM, 14
JSX026 J003, RTNCD 1012 REQSESS/TERMSESS OPEN OPNSEC FAILED
SENSE 08570002
```

Message IST891I may be issued with the IST663I message group and provides information about the identity of the nodes involved.

Probable cause

An incorrect CoS table was referenced. The NetView program also has a CoS table, and the NetView program library was concatenated in front of the VTAM library, causing the wrong table selection.

User response

To ensure that you are using the correct table, enter:

```
D NET,ID=resourcename
```

Check the library search order to ensure that there are no duplicate table names. Reassemble the table, and check the condition codes. If the condition code received is what you expected, relink it to the table.

Message IST718I and IST719I received when activating a CDRM

Problem statement

The messages IST718I and IST719I are received during the activation of a CDRM.

Common symptoms

The following messages are displayed:

```
IST718I ADDRESS INVALID FOR NETID=cdmnetid CDRM=cdmname CODE=X'code'  
IST719I SUBAREA subarea ELEMENT e1
```

Probable cause

The message is usually a symptom of a duplicate definition for a network address.

The duplicate may have been defined using the SUBAREA and ELEMENT parameters in another CDRM definition or in a GWPATH definition in a gateway NCP.

The duplicate may have been defined using the ADJNETSA and ADJNETEL parameters in another CDRM definition or in a GWPATH definition.

User response

If the duplicate network cannot be found by inspecting other definitions, run a VTAM internal trace with OPT=(NRM,MSG). When the trace is completed and IST718I and IST719I have been issued, use the console DUMP command to dump the VTAM region and CSA.

- Locate the MSG entry for the IST718I message in the trace.
- Before the message entry there should be an SRTF entry with a nonzero return code, usually 04. This SRTF entry points either directly or indirectly to the duplicate.
- See the SRTx VIT entry in z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT.
- The SRT entry address field points to an SRT entry that can be located in the dump of VTAM. This SRT entry plus X'10' points to the definition that has the duplicate network address.

Messages IST154I, IST561I, IST999E, IST566I, and IST930I received for APPL-APPL storage expansion failure

Problem statement

If APPL-APPL sessions are not paced at the session level, storage expansion failures can occur with messages IST154I, IST561I, IST999E, IST566I, and IST930I. The job entry subsystem (JES) has experienced this failure.

Common symptoms

The following messages are symptoms of storage expansion failures.

```
IST154I EXPANSION FAILED FOR LFBUF OR IOBUF BUFFER POOL  
IST561I STORAGE UNAVAILABLE  
IST999E VTAM MESSAGE LOST-INSUFFICIENT STORAGE  
IST566I STORAGE UNAVAILABLE xxxx SUBPOOL xxx  
IST930I LU-LU SESSION USING 15% OF IOBUF
```

Probable cause

If an APPL-APPL session is not paced at the session level, there is no limit to the number of VTAM I/O buffers that the session can use.

User response

Code VPACING operands on the APPL definition statements for both

applications, and code a nonzero value for the SSNDPAC parameter on the LOGMODE operand for the secondary LU. To verify pacing for the session, start a buffer trace with ID=APPLID specified before you start the APPL-APPL session. The BIND RU will contain the pacing values for the session.

Missing VTAM trace records

Problem statement

The expected output data is missing from a VTAM trace that was run with GTF active.

Common symptoms

There is no VTAM data, missing VTAM data, or unwanted data in the GTF trace data set.

Probable cause

When TRACE=USR is specified, GTF collects all USR events issued in the MVS system.

User response

To select the events you want to trace, specify USRP on the GTF macroinstruction and select the required event identifiers (EIDs) as shown in the following examples:

VTAM buffer EIDs: FEF FF1 FF0 (EFEF EFF1 EFF0)

VTAM line trace EIDS (not formatted by GTFTRACE): FE4 FF2 (EFE4 EFF2)

VTAM I/O trace EID: FE1 (EFE1)

VTAM internal trace EID: FE1 (EFE1)

See "Activating network traces" on page 308 for more information.

Note: To prompt the system for VTAM records, specify USRP in the parameter field of the GTF procedure. You must code a GTF procedure that is used by VTAM only. If you do not, you will get GTF USR output that contains unwanted records.

Partial input for TGET received by TSO

Problem statement

A TSO application program receives partial input for the TGET macroinstruction.

Common symptoms

The TSO application program does not receive the entire data-stream buffer from a device. A partial buffer from a device will cause the application to enter a wait state. If the host application program then issues a second TGET, the second section of the buffer is returned to the host application program before processing for the first TGET is completed.

Probable cause

The TSO application issued the set full-screen mode (STFSMODE) macroinstruction without specifying the NOEDIT option. The error occurs most often after the application program sends a read partition query (RPQ) to the device. Many newer devices return the attribute byte X'1E' that is returned in the RPQ entry. TSO interprets the X'1E' as an end-of-input field mark. The NOEDIT option of the STFSMODE macroinstruction prevents TSO VTAM from validating the input data. This causes the entire buffer to be returned to the application program.

User response

Verify the options on the STFSMODE macroinstruction. If STFSMODE is correct, see "Incorrect output problems" on page 112 for more information.

Sense code 08010000 received with CINIT failure

Problem statement

A CINIT request fails with the sense code 08010000 if an application rejects a terminal logon request by issuing the CLSDST macroinstruction.

Common symptoms

The following message is displayed:

```
IST663I request REQUEST FAILED, SENSE=08010000
```

The logon from a terminal fails with the USS message USSMSG07.

Probable cause

When an SLU logs on to an application, VTAM builds a CINIT RU and schedules the LOGON exit routine for the application PLU. If the application is not prepared to accept a session with this SLU, it rejects the logon by issuing a CLSDST macroinstruction. If the application does not supply sense code information about the CLSDST, VTAM builds a negative CINIT response with the sense code 08010000. In many cases, the application will also issue a message indicating the reason for the logon rejection.

User response

Check the message log for a message indicating a failure for this application. Run a buffer trace on the application name to see whether the CINIT passed to the application. The VTAM internal trace with the API option contains data about the LOGON exit and the CLSDST macroinstruction.

Sense code 083500xx or 08210000 received with BIND failure

Problem statement

BIND failure occurs with sense code 083500xx or 08210000.

Common symptoms

The following messages are displayed in response to a terminal logon request:

```
IST663I BIND REQUEST FAILED, SENSE=083500xx
or
IST663I BIND REQUEST FAILED, SENSE=08210000
and
USSMSG07 luname UNABLE TO ESTABLISH SESSION-BIND FAILED
WITH SENSE sense
```

Probable cause

The sense codes indicate that the BIND contains parameters that are not valid. The sense code 08210000 gives no further explanation. Sense code 083500xx supplies an index (xx) into the BIND that identifies the bytes that the BIND receiver cannot interpret.

VTAM extracts BIND parameters from the LOGMODE entry associated with the logon, based on the LOGON command, the USSPARM PARM=LOGMODE from the USSTAB, or the default on the LU definition specified by DLOGMOD. The source of the BIND parameters can also be the application, which may override many of the parameters supplied by VTAM when the OPNDST macroinstruction is issued. When the requested

LOGMODE cannot be found, VTAM may use a default LOGMODE (ISTCOSDF), which may contain session parameters that are unacceptable to the application.

User response

Run a buffer trace on the application name for a terminal session logon to an application that is rejected by the BIND with a sense code of 083500xx or 08210000. This traces the CINIT request, which includes the VTAM supplied parameters. If the BIND that follows the CINIT request does not match these parameters, they were changed by the application. The documentation for the rejecting LU should list its required BIND parameters.

To prevent VTAM from using the default LOGMODE (ISTCOSDF), ensure that the requested LOGMODE is defined in the specified LOGMODE table. See the *z/OS Communications Server: SNA Network Implementation Guide* and *z/OS Communications Server: SNA Resource Definition Reference* for more information on ISTCOSDF.

Sense code 800A0000 or no message, and sessions end unexpectedly

Problem statement

A session ended unexpectedly and either no message is received or an exception request (EXR) with a sense code of 800A0000 flows to the destination LU.

Common symptoms

Upon receiving the sense code 800A0000, the LU might return the code in a response. Some LUs will include the code in an UNBIND.

Probable cause

If a path information unit (PIU) is too large to be passed from one PU type 4 or type 5 to another, an exception request (EXR), containing sense code 800A0000 and up to 3 bytes of the RU, may be generated. See Table 48 on page 649 to determine what document describes the building of the EXR.

User response

For VTAM and NCP nodes in the session path, check the following definition values for each configuration used:

- VTAM to channel-attached NCP: VTAM will take the smaller of the following two values:
 - MAXDATA value on the PCCU definition statement (or on the LINE definition statement for a channel-attached NCP).
 - Value sent in the XID of the maximum PIU size for the NCP. This number will be the product of the BFRS value from the BUILD definition statement and the TRANSFER value from the channel adapter LINE definition statement for 3745 or 3720 with V5 NCP or from the BUILD definition statement for other NCPs.
- Channel-attached NCP to VTAM: The product of the MAXBFRU value from the HOST definition statement (or from the LINE definition statement for a channel-attached NCP) and the IOBUF size from the VTAM start options. This value will be the maximum size that can flow from the NCP to the host.
- VTAM-to-VTAM connection across a channel-to-channel interface:
 - If both VTAMs have the CTCA enhancement: The product of the MAXBFRU value from the CTCA LINE definition statement and the IOBUF size from the VTAM start options.

- If one or neither of the VTAMs has the CTCA enhancement: The product of the MAXBFRU value from the LINE macro and the IOBUF size from the VTAM start options.
- NCP to link-attached NCP: The product of the TRANSFR value on the LINE definition statement and the BFRS value on the BUILD definition statement.
- NCP to link-attached VTAM: The product of the MAXBFRU value from the CA LINE definition statement and the IOBUF size from the VTAM start options.

Note:

1. The definition statements for all PU type 4 or type 5 nodes on the session path must be checked, because any PU type 4 or type 5 can change the PIU into an 800A0000 exception request.
2. The information in this problem description is from information APAR II03990.

Sense codes 10030000 and 08090000 received when activating an NCP

Problem statement

Some resources fail to activate correctly when a new NCP is activated.

Common symptoms

VTAM commands return sense codes 10030000 and 08090000.

Probable cause

The resource resolution table (RRT) created when a new NCP was generated did not replace the previous RRT, and the system is still referring to the old RRT.

User response

When you generate a new NCP, either rename the new NCP or use another method to ensure that the old RRT is replaced with the new RRT.

Session failure with sense code 0888000x

Problem statement

An attempt to establish a session fails with sense code 0888000x in an intermediate VTAM along the session setup path.

Common symptoms

The session establishment is terminated.

Probable cause

The intermediate VTAM that set the 0888000x sense codes is operating with NQNMODE=NAME or is a VTAM version lower than V4 and therefore cannot define multiple resources with the same name, even if the network identifiers are different.

User response

Change the intermediate domain to operate with NQNMODE=NQNAME to allow definition of multiple resources with the same name and different network identifiers, or reroute the session through another path.

Common problems in APPN networks

Table 2 on page 23 includes a brief description of several common problems that occur in APPN networks. For additional information, go to the page indicated.

Table 2. Index of common problems in APPN networks

Problem	See page
Best path not taken for the session	"Session did not take the best path" on page 36
Message IST489I received during session takeover	"Message IST489I or IST1272I received during session takeover" on page 26
Message IST264I (required CoS entry undefined) and message IST663I (sense code 08610000)	"Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry" on page 27
Message IST663I (sense code 08610000) and message IST264I (required CoS entry undefined)	"Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry" on page 27
Messages IST1097I and IST1280I (sense code 08A00005) received with CP-CP session failure	"Messages IST1097I and IST1280I (sense code 08A00005) received with CP-CP session failure" on page 27
Messages IST1110I, IST1112I, IST1765I, and IST1766I received during CP-CP session activation failure	"Messages IST1110I, IST1112I, IST1765I, and IST1766I received during CP-CP session activation failure" on page 28
Messages IST1110I and IST1113I issued during CP-CP session activation failure	"Messages IST1110I and IST1113I received during CP-CP session activation failure" on page 28
Messages IST1110I and IST1246I issued during CP-CP session activation failure	"Messages IST1110I and IST1246I received during CP-CP session activation failure" on page 28
Messages IST1110I, IST1246I, and IST1280I (sense code 80050000) issued during CP-CP session activation failure	"Messages IST1110I, IST1246I, and IST1280I received during CP-CP session activation failure" on page 29
Messages IST1110I and IST1280I (sense code 08B50000) issued during CP-CP session activation failure	"Messages IST1110I and IST1280I (sense code of 08B50000) received during CP-CP session activation failure" on page 30
Messages IST1110I and IST1280I (sense code 08910006) issued during CP-CP session activation failure	"Messages IST1110I and IST1280I (sense code 08910006) received during CP-CP session activation failure" on page 30
Messages IST1110I, IST1280I (sense code 101E000A), and IST1356I issued during CP-CP session activation failure	"Messages IST1110I, IST1356I, and IST1280I (sense code 101E000A) received during CP-CP session activation failure" on page 29

Table 2. Index of common problems in APPN networks (continued)

Problem	See page
Message IST1272I received during session takeover	"Message IST489I or IST1272I received during session takeover" on page 26
Messages IST1774I and IST1775I received during LU-LU session activation	"Messages IST1774I and IST1775I received during LU-LU session activation" on page 30
MNPS recovery not successful	"MNPS session recovery error" on page 37
Resource not found but resource exists in network	"Resource not found but resource exists in network" on page 31
Sense code 08210002 issued during a session activation failure	"Sense code 08210002 received with session activation failure" on page 31
Sense code 0821000A issued during a session activation failure	"Sense code 0821000A received with session activation failure" on page 32
Sense code 083B0001 issued and session lost	"Sense Code 087D000A or 083B0001 received and session lost during takeover" on page 33
Sense code 08610000 (message IST663I) and message IST264I (required CoS entry undefined)	"Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry" on page 27
Sense code 087D0001 issued during a session activation failure	"Sense code 087D0001 received with session activation failure" on page 32
Sense code 087D000A issued and session lost	"Sense Code 087D000A or 083B0001 received and session lost during takeover" on page 33
Sense code 08910006 and messages IST1110I and IST1280I issued during CP-CP session activation failure	"Messages IST1110I and IST1280I (sense code 08910006) received during CP-CP session activation failure" on page 30
Sense code 08A00005 received unexpectedly with CP-CP session failure	"Messages IST1097I and IST1280I (sense code 08A00005) received with CP-CP session failure" on page 27
Sense code 08B50000 and messages IST1110I and IST1280I issued during CP-CP session activation failure	"Messages IST1110I and IST1280I (sense code of 08B50000) received during CP-CP session activation failure" on page 30

Table 2. Index of common problems in APPN networks (continued)

Problem	See page
Sense code 10145046 (AS/400®) issued during a session activation failure	“Sense code 10145046 received with session activation failure” on page 33
Sense code 101E000A and messages IST1110I, IST1356I, and IST1280I issued during CP-CP session activation failure	“Messages IST1110I, IST1356I, and IST1280I (sense code 101E000A) received during CP-CP session activation failure” on page 29
Sense code 80050000 and messages IST1110I, IST1246I, and IST1280I issued during CP-CP session activation failure	“Messages IST1110I, IST1246I, and IST1280I received during CP-CP session activation failure” on page 29
Sense code 80130000 issued during a session activation failure	“Sense code 80130000 received with session activation failure” on page 34
Sense code 80130104 and path problems	“Example: Solving path problems” on page 157
Sense code 80140001 issued during a session activation failure	“Sense code 80140001 received with session activation failure” on page 34
Sense code 80140002 issued during a session activation failure	“Sense code 80140002 received with session activation failure” on page 35
Sense code 80140005 issued during a session activation failure	“Sense code 80140005 received with session activation failure” on page 35
Session did not take the best path	“Session did not take the best path” on page 36
Session established with nonlocal instead of local application program	“Session established with nonlocal instead of local application program” on page 36
Session lost with sense code 083B0001	“Sense Code 087D000A or 083B0001 received and session lost during takeover” on page 33
Session lost with sense code 087D000A	“Sense Code 087D000A or 083B0001 received and session lost during takeover” on page 33
Storage problem	“Storage problem procedure” on page 97

Descriptions of common problems in APPN networks

This information includes examples of common problems in APPN networks. See Table 2 on page 23 for an index of these problems.

Message IST489I or IST1272I received during session takeover

Problem statement

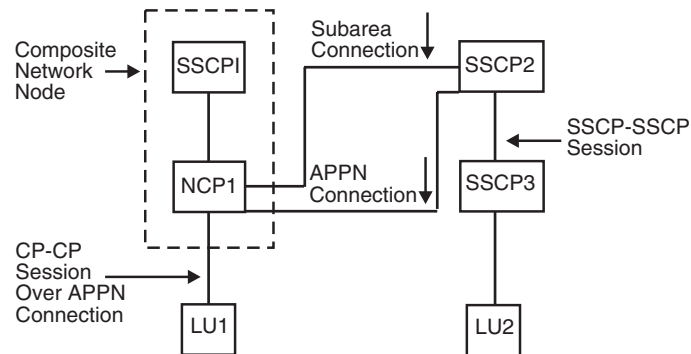
Message IST489I or IST1272I is received, indicating that VTAM cannot define a resource.

Common symptoms

Message IST489I or IST1272I is received for a resource during takeover processing. The resource can be a dependent LU.

Probable cause

A scenario similar to the following condition can cause this problem:



- Boundary function CP-CP sessions exist between SSCP1 and SSCP2. There is also an FID4 connection between SSCP2 and NCP1.
- There is an SSCP-SSCP session between SSCP2 and SSCP3.
- Dependent LU1, connected from NCP1, is owned by SSCP1.
- An LU-LU session is between LU1 and LU2. LU1 is known to SSCP2 as an APPN LU through the APPN connection between SSCP1 and SSCP2 through NCP1.
- The command VARY INACT TYPE=G was entered from SSCP1, which caused session takeover of the line from SSCP1 to SSCP2. SSCP2 owns LU2, the independent LU.

The failure, triggering message IST489I or message IST1272I, occurred because LU1 is known to the SSCP2 host as an independent APPN LU. SSCP2 cannot own the independent LU and the dependent LU at the same time.

User response

The independent LU is freed when the last session ends. After this occurs, you must activate the dependent LU before it can be enabled.

Message IST663I (sense code 08610000) and IST264I received for undefined CoS entry

Problem statement

A required CoS entry is UNDEFINED.

Common symptoms

The following messages are received:

```
IST663I  request REQUEST FAILED, SENSE=08610000
IST264I  REQUIRED COS luname UNDEFINED
HASP208  LOSTTERM SCHEDULED SNA, VTAM, 14
JSX026   J003, RTNCD 1012 REQSESS/TERMESS OPEN OPNSEC FAILED
          SENSE 08570002
```

Message IST891I may be issued with the IST663I message group and provides information about the identity of the nodes involved.

Probable cause

An incorrect CoS table is referenced. The NetView program also has a CoS table, and the NetView program library was concatenated in front of the VTAM library, causing the wrong table selection.

User response mode-to-Class-of-Service (CoS)

To ensure that you are using the correct table, enter:

```
D NET, ID=resourceName
```

Check the library search order to ensure that there are no duplicate table names. Reassemble the table, and check the condition codes. If the condition code received is what you expected, relink it to the table.

Messages IST1097I and IST1280I (sense code 08A00005) received with CP-CP session failure

Problem statement

CP-CP session failure occurs with sense code 08A00005.

Common symptoms

The following message group is received:

```
IST1097I  CP-CP SESSION WITH partner_cpname TERMINATED
IST1280I  SESSION TYPE = CONLOSER - SENSE = 08A00005
IST314II  END
```

Probable cause:

If CP-CP sessions have been deactivated with this sense code, it is likely that the topology database update (TDU) flowing between the two nodes has been lost because of a storage depletion condition on either the sending or receiving end of the TDU flow.

User response

If the CP-CP sessions do not come backup automatically, enter:

```
V ACT, ID=partner_cpname, IDTYPE=CP
```

If VTAM is experiencing temporary storage allocation problems, you might want to wait for the condition to clear before attempting to restart the session.

When the CP-CP session is restarted, TDUs will be exchanged so that the missing information in the lost flow will be recovered.

Messages IST1110I, IST1112I, IST1765I, and IST1766I received during CP-CP session activation failure

Problem statement

End node operator's attempt to activate CP-CP session pair by activating adjacent network node fails and messages IST1110I, IST1112I, IST1765I, IST1766I, and IST314I are issued at the end node.

Common symptoms

Activation of new CP-CP session pair with adjacent CP specified in message IST1110I is terminated. Messages IST1110I, IST1112I, IST1765I, IST1766I, and IST314I are displayed.

Probable cause

While attempting to activate a contention winner CP-CP session with the network node specified in message IST1110I, the end node determined that it already had an active CP-CP session pair with a different network node. If an end node already has a network node server, it does not accept a new CP-CP session with another network node.

User response

Before attempting to activate a CP-CP session pair between an end node and a network node, enter the `D NET,NETSRVR,SCOPE=ONLY` command at the end node to verify that no CP-CP sessions with a network node exist.

Messages IST1110I and IST1113I received during CP-CP session activation failure

Problem statement

Attempt by operator of end node to activate CP-CP session pair with adjacent network node by activating adjacent network node fails with issuance of messages IST1110I, IST1113I, and IST314I.

Common symptoms

Activation of new CP-CP session pair with adjacent CP specified in message IST1110I is terminated. Messages IST1110I, IST1113I, and IST314I are displayed at the end node.

Probable cause

The operator entered from an end node a `V NET,ACT,ID=cpname` command, where `cpname` is also an end node. CP-CP sessions are not permitted between end nodes.

User response

Make sure that the start lists for the two nodes do not both specify end node as the type of node being started.

Messages IST1110I and IST1246I received during CP-CP session activation failure

Problem statement

Attempt to activate CP-CP session pair between end node and network node fails after activation of CP-capable link.

Common symptoms

Activation of new CP-CP session pair with network node specified in message IST1110I is terminated. Message IST1110I is displayed along with message IST1246I at the end node.

Probable cause

While in the process of bringing up the contention winner CP-CP session,

the end node determined that the network node named in IST1110I is not explicitly named in the end node's network node server list and that there is no nameless entry in the network node server list.

User response

Either perform an operator activation of the CP-CP session by entering V NET,ACT,ID=adjacent_cpname at the end node or modify the network node server list to include either an explicit entry for the required network node or a nameless entry.

Messages IST1110I, IST1246I, and IST1280I received during CP-CP session activation failure

Problem statement

Attempt to activate CP-CP session pair between end node and network node fails.

Common symptoms

Activation of new CP-CP session pair with network node specified in message IST1110I is terminated. Message IST1110I is displayed along with message IST1246I at the end node. Message IST1280I displays a sense code of 80050000.

Probable cause

While in the process of bringing up the conloser CP-CP session, the end node determined that the network node named in IST1110I is not explicitly named in the end node's network node server list and that there is no nameless entry in the network node server list.

User response

Either perform an operator-activation of the CP-CP session by entering V NET,ACT,ID=adjacent_cp_name at the end node or modify the network node server list to include either an explicit entry for the required network node or a nameless entry.

Messages IST1110I, IST1356I, and IST1280I (sense code 101E000A) received during CP-CP session activation failure

Problem statement

Attempt to activate CP-CP session pair between end node and network node fails.

Common symptoms

Activation of new CP-CP session pair with network node specified in message IST1110I is terminated. The operator at the end node sees messages IST1110I, IST1356I, IST1280I, and IST314I. IST1280I displays the sense code 101E000A.

Probable cause

The end node's network node server list entry for the network node failed to specify SLUINIT=OPT or, in the absence of an explicit entry for that node, the nameless entry failed to specify SLUINIT=OPT. The network node is probably an AS/400, NS/2, or Personal System/2 computer, none of which provides the network node server capabilities provided by VTAM network nodes. CP-CP sessions between VTAM end nodes and such network nodes are allowed only if the end node's network node server list specifies SLUINIT=OPT.

User response

Modify the network node server list to specify SLUINIT=OPT on either the

explicit entry for the required network node server or on the nameless entry. Activate the modified network node server list definition deck, and then reactivate the session.

Messages IST1110I and IST1280I (sense code 08910006) received during CP-CP session activation failure

Problem statement

Attempt to activate CP-CP session pair between network nodes in two different networks fails.

Common symptoms

Activation of new CP-CP session pair between this network node and adjacent network node specified in message IST1110I is terminated.

Probable cause

A CP-CP session pair is not permitted between network nodes located in different networks unless you have specified BN=YES to enable the VTAM border node function. The messages indicate that CP-CP sessions were attempted between two network nodes in different networks.

User response

If you want a nonnative relationship, ensure that BN=YES is coded to enable border node support. Also, ensure that NATIVE=YES is not coded on a PU or ADJCP statement that represents the partner node. If you want a native relationship, modify the VTAM start lists for the specific nodes so that both start lists specify the same network.

Messages IST1110I and IST1280I (sense code of 08B50000) received during CP-CP session activation failure

Problem statement

Attempt to activate CP-CP session pair between end node and network node fails with sense code 08B50000, as indicated by message IST1280I.

Common symptoms

Activation of new CP-CP session pair with adjacent CP specified in message IST1110I is terminated. Message IST1110I is displayed along with message IST1280I, which displays a sense code of 08B50000.

Probable cause

The sense code indicates that the end node bringing up the contention-loser session does not require a CP-CP session pair with the network node specified in message IST1110I. The end node determined that it has an active CP-CP session with a different network node. If an end node already has a server, it will not accept a new CP-CP session with another network node.

User response

Before attempting to activate a CP-CP session pair between an end node and a network node, enter the D NET,NETSRVR,SCOPE=ONLY command at the end node to verify that no CP-CP sessions with a network node already exist.

Messages IST1774I and IST1775I received during LU-LU session activation

Problem statement

LU-LU session activation completes successfully with the issuance of messages IST1774I, IST1775I, IST664I, IST889I, and IST314I at the composite network node (CNN).

Common symptoms

The following messages are received:

```
IST1774I OPTIMAL CNN ROUTE NOT CHOSEN - ENTRY/EXIT SUBAREA MISMATCH
IST1775I CNN ENTRY SUBAREA = subarea   CNN EXIT SUBAREA = subarea
IST664I REAL  OLU=1uname                REAL  DLU=1uname
IST889I SID = sessid
IST314I END
```

Probable cause

An optimal CNN route exists and was not chosen during session activation. Non-optimal routes might result because the topology of the CNN is not known by the APPN topology and route selection process, or because the route was calculated by a non-VTAM node that does not support the use of subarea numbers in route calculation.

User response

Change the APPN TG characteristics. For more information about APPN TG characteristics, See the z/OS Communications Server: SNA Network Implementation Guide.

To suppress message group IST1774I, CNNRTMSG=SUPPRESS can be specified as the start option value or modified with the MODIFY VTAMOPTS command.

Resource not found but resource exists in network

Problem statement

A resource exists in the network but is not found by a search.

Common symptoms

The directory services management exit routine either rejects or limits the search scope.

Probable cause

The resource was not registered to its network node server.

User response

Register the resource to the network node server. For more information on the directory services management exit routine, see z/OS Communications Server: SNA Customization.

Sense code 08210002 received with session activation failure

Problem statement

Session activation failed with the sense code 08210002.

Common symptoms

An attempt to establish a session failed with the sense code 08210002 (mode name not valid).

Probable cause

The sense code indicates that the logon mode name associated with the session request was not found in the table or in the default logon mode table (ISTINCLM).

User response

Verify that the requested logon mode name is defined as follows:

- In a subarea-only environment, the mode name must be defined in the SSCP associated with the SLU.
- In an APPN-only environment, the mode name must be defined in the origin and destination nodes, as well as the origin and destination node servers if the origin or the destination is owned by an end node.

- In a combined APPN and subarea environment, the mode name must be defined at the APPN node that owns the origin or destination, at the node server if it is an end node, and at the interchange nodes that represent the subarea entry point. If the SLU is owned by a subarea node other than the interchange node representing the subarea entry point, the mode name must also be defined on the owning subarea.

Note: Because mode table names are not carried on APPN line flows, a user-defined mode table is used only at the SSCP for the SLU. Other nodes defining the mode must define the SLU in the default logon mode table.

Refer to z/OS Communications Server: SNA Network Implementation Guide for details on mode to CoS resolution in an APPN or in a combined subarea and APPN environment.

Sense code 0821000A received with session activation failure

Problem statement

Session activation failed with the sense code 0821000A.

Common symptoms

An attempt to establish a session failed and the sense code 0821000A (mode table not found) was returned.

Probable cause

The sense code indicates that the mode table associated with the LU was not found.

User response

Verify that the specified table exists, and activate it.

Sense code 087D0001 received with session activation failure

Problem statement

Session activation failed with the sense code 087D0001.

Common symptoms

An attempt to establish a session failed and the sense code 087D0001 (routing exhausted) was returned. Messages IST894 and IST895 indicate that one of the adjacent SSCPs tried is ISTAPNCP with a failure sense code of 087F0001 (resubmit requested for a request that was already resubmitted).

Probable cause

Possibly one of the following conditions:

- If messages IST894I and IST895I are issued, one of the adjacent SSCPs was ISTAPNCP with a failure sense code of 087F0001. This indicates that VTAM knows which node owns the LU but is not able to route a directed search to that node to verify the availability of the LU.
- There is no SSCP-SSCP session.
- The half-session control block (HSCB) count is too low in the NCP to handle the number of sessions. A possible solution to this problem is to code a larger value on the ADDSESS keyword of the BUILD definition statement and regen.
- Both sides are using the same SSCP name.

User response

Verify that a valid search path exists. This includes CP-CP sessions, a subarea path, or both. One possible source of the problem is the absence of

a CP-CP session between two nodes that share an active link that is CP-CP capable. If this situation occurs, take one of the following actions:

- Reactivate the CP-CP session.
- Deactivate the link, and reactivate it as a link that is not CP-CP capable. This notifies topology and routing services that the link is no longer available for use in directed search routing.

Sense Code 087D000A or 083B0001 received and session lost during takeover

Problem statement

Session lost during takeover with sense code 087D000A or 083B0001.

Common symptoms

An attempt to take over a switched connection that is defined with ANS=CONTINUE results in a session or sessions being lost. A message states that a BFSESSINFO request failed with the sense code 087D000A (routed through same SSCP twice) or with the sense code 083B0001 (duplicate PCID).

Probable cause

The problem might be that a connection-network-capable control point (CP) on the connection network does not have a complete system definition.

User response

If you have a connection network, check the resource definitions on each CP connected to the network. Any connection-network-capable CP must define both of the following connections:

- Its own connection to the connection network
- Connections to any CPs on the connection network that are not connection-network-capable

Sense code 10145046 received with session activation failure

Problem statement

Session activation failed with the AS/400 sense code 10145046.

Common symptoms

An attempt was made to establish a session from one AS/400 to another AS/400 across a VTAM network. The session failed to complete, and route selection errors occurred on the initiating AS/400.

Probable cause



NETC

NETA

NETB

The AS/400 (NETC) is sending VTAM (NETA) a CDINIT that specifies a session in a third network that is not supported in this release of VTAM. VTAM rejects the session, and the AS/400 returns the sense code 10145046 to VTAM.

User response

Verify that APPN sessions across three networks are not present in your

system. Sessions across three networks using APPN links are not supported by VTAM unless you have specified BN=YES to enable the VTAM border node function.

Sense code 80130000 received with session activation failure

Problem statement

Session activation failed with the sense code 80130000.

Common symptoms

An attempt to establish a session failed, and the sense code 80130000 (Class of Service not available) was returned.

Probable cause

This sense code indicates that the subarea Class of Service (CoS) is not known. (In contrast, sense code 80140002 is issued by topology and routing services and indicates that the APPN CoS is not known.)

User response

Verify that the node issuing the sense code has a usable subarea Class of Service for the mode associated with the session request. This node is usually the primary logical unit (PLU) host, an intermediate gateway VTAM, or a gateway VTAM.

Message IST891I may be issued and provides the name of the failing node. See the description of message IST891I in z/OS Communications Server: SNA Messages for additional information.

This problem can occur when a mode table is copied from one node to another, and the subarea classes of service specified by the table no longer map to valid CoS names defined at that node.

Sense code 80140001 received with session activation failure

Problem statement

Session activation failed with the sense code 80140001.

Common symptoms

An attempt to establish a session failed, and the sense code 80140001 (no route exists) was returned. This sense code indicates that no APPN route was found from the origin node to the destination node that meets the requirements of the requested Class of Service.

This can be due to any of several possible causes.

Probable cause 1

There is not an active APPN route between the origin and destination.

User response 1

Examine your network configuration to determine whether a valid path does exist. Use the DISPLAY TOPO command to verify that the topology database currently shows the links in the path as operational.

Probable cause 2

Although a valid APPN path exists, the characteristics of the nodes and links in the operational paths do not meet the requirements of the specified Class of Service.

User response 2

Check the following items:

1. Verify that the mode name specified on the request maps to the intended Class of Service.

2. Examine the LINEROW and NODEROW operands in the Class of Service definition to determine the allowable ranges for the link and node characteristics.
3. Use the DISPLAY TOPO command to view the characteristics of the nodes and TGs in the likely paths. Look for problems such as:
 - a. Nodes in the path are congested or have route resistance values outside the limits set by the Class of Service.
 - b. The CoS definition required secure links, but no path exists consisting exclusively of secure TGs.
 - c. High capacity (speed) was required by the CoS definition, but no path exists in which all of the links are fast enough to meet the specified minimum capacity.

Probable cause 3

The destination exists in a subarea network, or in another APPN network that is accessed through a subarea network, but paths acceptable to the specified Class of Service definition do not exist to all interchange nodes representing subarea entry points that can be used to reach the destination.

User response 3:

As specified in z/OS Communications Server: SNA Network Implementation Guide, if a destination can be reached by exiting one APPN network through two or more different interchange nodes, paths acceptable for the Class of Service to be used must be available to all of those possible exit interchange nodes. Verify this by examining the CoS definition and the characteristics of the paths to the possible exit interchange nodes.

Sense code 80140002 received with session activation failure

Problem statement

Session activation failed with the sense code 80140002.

Common symptoms

An attempt to establish a session failed, and the sense code 80140002 (not valid APPN CoS name received) was returned.

Probable cause

The sense code indicates that the APPN CoS definition was not found. The definition might not exist at a node that is performing mode-to-CoS resolution, or the mode-to-CoS mapping specified in the mode table might not be mapping to the intended CoS name.

User response

Examine the mode definition to determine the APPN CoS name. Verify that this definition exists in the VTAMLST members at the nodes that resolve the mode to an APPN Class of Service. Activate the member to be sure that the definition is active. If APPN CoS substitution is enabled (by specifying the APPN CoS start option), verify that the CoS it specifies has been activated.

Sense code 80140005 received with session activation failure

Problem statement

Session activation failed with the sense code 80140005.

Common symptoms

An attempt to establish a session failed, and the sense code 80140005 (RSCV exceeds the maximum length) was returned.

Probable cause

The sense code indicates that the number of hops between the origin and destination nodes was too large, so the attempt to build the Route Selection control vector failed.

User response

Examine your network configuration to determine how many hops would be expected in the best route for the requested APPN Class of Service. If the number of hops within a single APPN network is greater than six, you may need to provide a more direct origin to the destination path.

Session established with nonlocal instead of local application program**Problem statement**

The session was intended to be established with a local application program; however, it was established with a nonlocal application program.

Probable cause

The local application program is not yet active, and a local application program served by another node is registered to the nonlocal application program network node server.

User response

Be sure that the local application program is active before you attempt to log on to it.

Session did not take the best path**Problem statement**

A session took one of the following paths:

- The session took a path through the subarea network even though a better path existed through the APPN network, or the session took a path through the APPN network even though a better path existed through the subarea network.
- After a rapid transport protocol (RTP) connection switches to a new path, a session takes a path that requires it to visit the same node twice. For example, an LU-LU session between HOSTA and HOSTC goes from HOSTA through HOSTC to HOSTB and back to HOSTC.

Probable cause

- In the first situation described above, the SORDER operand or SSEARCH operand is coded with a value that indicates that the subarea network is to be searched before the APPN network is searched or that the APPN network is to be searched before the subarea network is searched.
- In the second situation described, HOSTA and HOSTB support rapid-transport protocol (RTP), but HOSTC supports only automatic network routing (ANR). Because one of the LUs resides on an ANR node (which cannot be the endpoint of an RTP connection), at least one hop of the session is not using high performance routing (HPR). During an RTP path switch, the non-HPR hops cannot change, but the RTP hops can. In some configurations, it is possible that the new path for the RTP connection will traverse some of the same nodes as the non-HPR portion of the original session route.

User response

- In the first situation, adjust the values on the SORDER and SSEARCH operands to suit your network.

- In the second situation, the session should continue, so no user action is necessary. If a temporary loss of connectivity forced RTP to switch paths, after restoring the connectivity, you can use the MODIFY RTP command to force VTAM to search for a better route for the RTP.

MNPS session recovery error

Problem statement

One or more MNPS sessions were terminated unexpectedly and were not recovered. This can be due to any of several possible causes.

Probable cause 1

Pathswitcher time set by HPRPST expired on other end of HPR pipe.

User response 1

Restart the application using automatic restart manager (ARM). Set HPRPST to allow more time (if possible) or recover more quickly the next time.

Probable cause 2

HPR connectivity is not consistent throughout the sysplex.

User response 2

Update network definitions to make sure that connectivity is consistent among all the MNPS nodes.

Probable cause 3

The VTAM is not in a sysplex, or the VTAM nodes are not in the same subplex, if subplexing is being used. Subplexing is being used if the XCFGRPID start option has been used to specify an XCF group ID suffix.

User response 3

Implement the recovery in a sysplex environment. If subplexing is being used, ensure that all nodes involved in MNPS sessions and session recovery are in the same subplex. VTAM nodes are in the same subplex if each node specifies the same 2-digit value on the XCFGRPID start option.

Probable cause 4

VTAM is not connected to the MNPS structure.

User response 4

Check the value of the STRMNPS start option. To determine the status of VTAMs connection, issue the command:

```
D NET,STATS,type=CFS,ID=MNPSstructurename
```

Probable cause 5

Pathswitcher time set by PSTIMER has expired.

User response 5

Restart the application using automatic restart manager (ARM). You can set the timer to allow more time using the application start definitions or else recover more quickly the next time.

Probable cause 6

The PERSIST=MULTI operand is not defined for the application. This sets the application for MNPS; if it is not defined, MNPS is not allowed.

User response 6

Terminate the application. Define PERSIST=MULTI on the APPL definition statement. Restart the application.

Probable cause 7

Session traverses a subarea path not on an RTP connection.

User response 7

Change the network configuration to ensure that the session is established over valid network routes. Ensure that SORDER is set to search APPN first.

Common problems in HPR networks

Table 3 includes a brief description of several common problems that occur in HPR networks. For additional information, go to the page indicated.

Table 3. Index of common problems in HPR networks

Problem	See page
LU-LU session initiation fails	"LU-LU session initiation fails"
LU-LU session initiation does not complete	"LU-LU session initiation does not complete" on page 39
No RTP connection established for CP-CP session	"No RTP connection established for CP-CP session pair" on page 39
LU-LU session established through ISR routing	"LU-LU session established using ISR routing" on page 39
RTP connection for LU-LU session does not include entire session	"RTP connection for LU-LU session does not include entire session path" on page 40
RTP connection experiences a path switch	"RTP connection experiences a path switch" on page 40
RTP path switch fails	"RTP path switch fails" on page 40
deactivate or MODIFY RTP for Route_Setup RTP ALS fails	"Deactivate or MODIFY RTP for Route_Setup RTP ALS fails" on page 41
Sense code 08770026	"Sense code 08770026 received on dial-out or dial-in for Enterprise Extender" on page 41
Sense code FFC80004	"Sense code FFC80004 received with dial-out for Enterprise Extender" on page 42

Descriptions of common problems in HPR networks

This information includes examples of common problems in HPR networks. See Table 3 for an index of these problems.

LU-LU session initiation fails

Problem statement

A severe error that prevents session initiation has occurred.

Common symptoms

Message IST663I is issued and contains a sense code describing the reason for the session initiation failure.

Probable cause

The RTP connection manager (RCM) experienced a severe error.

User response

Determine which node in the session path has experienced the error; obtain a VTAM dump, including a VTAM internal trace (VIT) with the HPR options.

LU-LU session initiation does not complete**Problem statement**

The LU-LU session activation begins but does not complete, and the session status is PRTPSTR.

Common symptoms

The LU-LU session activation begins, but message IST874I indicates that the session status is PRTPSTR.

Probable cause

The RTP connection between the session partner nodes has not fully activated.

User response

Specify the IOPURGE VTAM start option to clear waiting signals.

No RTP connection established for CP-CP session pair**Problem statement**

The CP-CP session pair is activated, but no RTP connections are established.

Common symptoms

No message IST1488I received before CP-CP session activation.

Probable cause

Connection between CP-CP capable nodes does not support the HPR Control Flows tower.

User response

Reconfigure the network to connect CP-CP capable adjacent nodes with resources supporting the HPR Control Flows tower.

LU-LU session established using ISR routing**Problem statement**

The LU-LU session activates, but no RTP connections are established.

Common symptoms

No message IST1488I was received before LU-LU session activation.

Probable cause

VTAM determined that the session path is not HPR capable:

- One or more connections between nodes in the session path are not HPR capable.
- The session path does not terminate with an HPR tower node. (If the terminating node is VTAM, it must specify the start option HPR=RTP or take the default).

User response

Reconfigure the session path to include HPR capable links and terminate with an HPR tower node.

RTP connection for LU-LU session does not include entire session path**Problem statement**

The LU-LU session activates and an RTP connection is established but includes only part of the session path.

Common symptoms

Message IST1487I contains a destination CPNAME that is not the name of the session partner node.

Probable cause

VTAM determined that part of the session path is not HPR capable:

- One or more connections between nodes in the session path are not HPR capable.
- The session path does not terminate with an HPR tower node. (If the terminating node is VTAM, it must specify the start option HPR=RTP or take the default).
- One or more nodes along the path were unable to perform HPR tower function.
- The session path enters the subarea and the HPR capable portion of the path ends in an interchange node at the subarea boundary.

User response

Reconfigure the session path to include HPR capable links and terminate with an HPR tower node. Ensure that all nodes in the path are able to perform HPR tower function.

RTP connection experiences a path switch**Problem statement**

A path switch operation is begun for an RTP connection.

Common symptoms

Message IST1494I indicates that path switch has been started for the RTP.

Probable cause

A resource in the RTP path has become inoperative.

User response

Determine which resource in the RTP path has become inoperative and restore that resource to operational status.

RTP path switch fails**Problem statement**

Message IST1494I indicates RTP path switch starts. Message IST1494I is reissued, indicating RTP path switch has failed.

Common symptoms

Message IST1495I indicates that no alternate route is available.

Probable cause

No alternate HPR route exists between the RTP edge nodes.

User response

Reconfigure the network to include an alternate HPR route.

Deactivate or MODIFY RTP for Route_Setup RTP ALS fails

Problem statement

A command has been issued for an RTP resource that has either an invalid node type or an invalid state for the command.

Common symptoms

Message IST607I indicates the command that failed for the specified RTP resource.

Probable cause

The command specified in message IST607I is not applicable for the RTP resource specified in message IST607I because the node type or state of the RTP resource is invalid for the operation that was requested.

User response

Issue a DISPLAY command for the RTP major node (ISTRTPMN) to verify RTP resource types and states. Reenter the command for a resource that is either the valid node type or in the valid state for the command.

Sense code 08770026 received on dial-out or dial-in for Enterprise Extender

Problem statement

Dial-out or dial-in failed with a sense code of 08770026.

Common symptoms

An attempt to dial out or dial in to establish a session failed and the sense code 08770026 was received.

Probable cause

The sense code indicates that the link station selected does not have HPR=RTP capability.

User response

Specify one of the following options:

- Start option HPR=RTP.
- Start option HPR=(RTP,ANR). Specify HPR=YES either on the PU or by the operation command activating the PU.

Sense code 1016000B received on dial-in for Enterprise Extender

Problem statement

Dial-in failed with a sense code of 1016000B.

Common symptoms

An attempt to dial-in or establish a session failed, and the sense code 1016000B was returned. Message IST1085I was issued on the host that rejected the dial-in.

Probable cause

The sense code indicates that a connection through TCP/IP has been established with identical TG number and CP name values. A duplicate CP name might be in the network.

User response

From the host where the 1016000B sense code was received, issue a DISPLAY EE,CPNAME= command, where CPNAME specifies the name of resource from the IST1085I message. Information about the CP with the active EE connection is displayed.

Sense code FFC80004 received with dial-out for Enterprise Extender

Problem statement

Dial-out failed with a sense code of FFC80004.

Common symptoms

An attempt to dial out or establish a session failed and the sense code FFC80004 was returned.

Probable cause

The sense code indicates that a connection through TCP/IP has already been established with identical local SAP, remote SAP, and IP address values.

User response

Verify that the remote SAP values specified on the PATH statements within the switched major nodes are unique.

Note: For a dial-through-a-connection network, the remote SAP value is the local SAP value of the node that is being dialed.

Common symptoms and associated VTAM problem types

If your problem was not described in “Common problems in subarea networks” on page 4 or “Common problems in APPN networks” on page 22, find the symptom you are experiencing in Table 4. The symptoms are listed alphabetically. Match your symptom to the appropriate VTAM problem type and go to the page indicated.

Table 4. Index of problem symptoms and associated VTAM problem types

Symptom	Problem type	See page
Abend message.	Abend	“Abnormal end (abend)” on page 57
Activating network nodes takes too long.	Performance	“Performance problem” on page 94
Application program cannot terminate.	Wait	“Wait” on page 61
Application programs and terminals cannot communicate.	Wait or Loop	“Wait” on page 61, “Loop” on page 81
Application program reports an unexpected return or sense code.	Incorrect Output or Message	“Incorrect output” on page 90, “Message problem” on page 87
Batch application program fails to complete.	Wait	“Wait” on page 61
Document is missing information or has wrong or ambiguous information.	Documentation	“Documentation problem” on page 100
Documents contradict each other.	Documentation	“Documentation problem” on page 100
Command is not completed.	Wait or Incorrect Output	“Wait” on page 61, “Incorrect output” on page 90
Commands cannot be entered on system console.	Loop	“Loop” on page 81
Commands take too long to complete.	Performance	“Performance problem” on page 94

Table 4. Index of problem symptoms and associated VTAM problem types (continued)

Symptom	Problem type	See page
Cursor is in the wrong position. This is probably an application program or VTAM definition error, such as using an incorrect logmode definition.	Incorrect Output	"Incorrect output" on page 90
Deactivating network nodes takes too long.	Performance	"Performance problem" on page 94
Error message.	Message	"Message problem" on page 87
Hung session, LU, or terminal.	Incorrect Output	"Incorrect output" on page 90
Hung system.	Wait	"Wait" on page 61
IKT error message.	Message	"Message problem" on page 87
IKT message is wrong or formatted improperly.	Message or Incorrect Output	"Message problem" on page 87, "Incorrect output" on page 90
IST error message.	Message	"Message problem" on page 87
IST message is wrong or formatted improperly.	Message or Incorrect Output	"Message problem" on page 87, "Incorrect output" on page 90
Keyboard locks unexpectedly.	Incorrect Output	"Incorrect output" on page 90
LOGON takes too long to complete.	Performance	"Performance problem" on page 94
LOGREC entries indicate an abend.	Abend	"Abnormal end (abend)" on page 57
LOGREC fills with repeated entries.	Loop or hardware	"Loop" on page 81
Message is wrong or formatted incorrectly.	Message or Incorrect Output	"Message problem" on page 87, "Incorrect output" on page 90
Message from application program.	Incorrect Output or Message	"Incorrect output" on page 90, "Message problem" on page 87
Message is sent to the wrong console.	Incorrect Output	"Incorrect output" on page 90
Message repeats continuously.	Loop	"Loop" on page 81
Message text does not explain a condition.	Message	"Message problem" on page 87
Message is missing text.	Message or Incorrect Output	"Message problem" on page 87, "Incorrect output" on page 90
Output data is formatted incorrectly. This is probably an application program or VTAM definition error, such as using an incorrect logmode definition.	Incorrect Output	"Incorrect output" on page 90
Path problem.	Performance	"Example: Solving path problems" on page 157
Performance is degraded after a network outage.	Performance	"Performance problem" on page 94

Table 4. Index of problem symptoms and associated VTAM problem types (continued)

Symptom	Problem type	See page
Printers stop.	Loop	"Loop" on page 81
PSWs point to a VTAM address.	Loop	"Loop" on page 81
Response time is slow.	Performance	"Performance problem" on page 94
Routing information is wrong.	Incorrect Output	"Incorrect output" on page 90
Storage message IST154I, IST562I, or IST561I-IST833I.	Storage	"Procedure steps" on page 97
System functions stop.	Loop	"Loop" on page 81
System light is on; Wait light is off.	Loop	"Loop" on page 81
Tapes stop.	Loop	"Loop" on page 81
Terminal user cannot log on, enter data, or log off.	Incorrect Output	"Incorrect output" on page 90
Terminal user gets unexpected response. This is probably an application program or VTAM definition error, such as using an incorrect logmode definition.	Incorrect Output	"Incorrect output" on page 90
Terminal user reports incorrect or missing data. This is probably an application program or VTAM definition error, such as using an incorrect logmode definition.	Incorrect Output	"Incorrect output" on page 90
Traffic ceases through a network component (BSC link, SDLC link, communication controller, control unit).	Wait	"Wait" on page 61
VTAM does not work as described in a document.	Documentation	"Documentation problem" on page 100
VTAM is not communicating with system console.	Wait or Loop	"Wait" on page 61, "Loop" on page 81
VTAM process issues an error message.	Message	"Message problem" on page 87

VTAM internal trace (VIT) analysis tool problems

This information describes how to diagnose problems that might occur while running the VIT analysis tool and includes the following topics:

- "Checklist for isolating the problem"
- "Common symptoms and actions" on page 45
- "Documenting an APAR for VIT analysis tool problems" on page 47

If you are having problems during installation of the tool, see the z/OS Communications Server: New Function Summary for additional information.

Checklist for isolating the problem

1. Does the problem exist with an ISPF panel or with the VIT analysis tool? The ISPF panel process creates the parameter data set, which is the input to the VIT analysis tool.
2. Are any errors noted in the JCL output? IBM publishes sample JCL, as described under "Step 2. Set up to run the tool" on page 366. However, you can change it to suit your environment. The JCL output indicates whether the tools entry

module, ISTRAFF1, is found. The JCL output contains an error message if the SUMMARY data set cannot be written to (for example, if the wrong DCB information was specified for the SUMMARY data set). Also check the JCL to verify that the right trace tapes or DASD data sets are specified and that multiple tapes are specified in the correct order.

3. Check the SUMMARY data set to determine whether any errors were noticed, such as the wrong DCB information being supplied for the LOG data set. The SUMMARY data set shows the input parameters and defaults used. The SUMMARY data set reports time stamps and types of records and VIT entry occurrences found on the trace. The SUMMARY data set should always be created, unless there was a problem with the SUMMARY data set itself, in which case the JCL or REXX output shows what happened.
4. Check the input parameters in the parameter data set with the syntax diagrams. The parameter data set is created by either the ISPF panel process or by coding the parameters directly using an editor. For more information on checking syntax diagrams, see Chapter 8, "Using the VIT analysis tool," on page 365.
5. Check the following output data sets:
 - LOG

The LOG data set shows whether counters overflowed, whether the trace wrapped, and so on. The LOG data set is always used unless an unrecoverable error prevents the tool from initializing completely. The LOG data set might contain only the title line and description (if a description exists).
 - VITEXT

The VITEXT data set is used only if the VIT extraction function is chosen. It is not used for storage analysis or RU-counting.
 - DETAILS

The DETAILS data set is used only if the storage analysis or RU-counting function is chosen. It is not used for VIT extraction.
 - OUTSTAN

The OUTSTAN data set is used only if storage-analysis-counting function is chosen and only if the outstanding option (to list unmatched allocate entries) is chosen. It is not used for RU counting or VIT extraction.

Common symptoms and actions

Use Table 5 to diagnose and correct problems.

Table 5. VIT analysis tool problems: Common symptoms and actions

Symptom	Action
Runs too long	<p>If the tool is taking a long time to run (several hours):</p> <ul style="list-style-type: none"> • If the MATCH option was specified for storage analysis: <ul style="list-style-type: none"> – Remove the MATCH option. – Specify only the few pools you are interested in when using the MATCH option. • Check to see if the tool is waiting on a tape to be mounted or to access a data set in use by some other job.

Table 5. VIT analysis tool problems: Common symptoms and actions (continued)

Symptom	Action
No output	<p>If no output is displayed (for example, no matching VIT entries are found, no RUs are counted, no storage VIT entries are found):</p> <ul style="list-style-type: none"> • Verify that the trace has VIT records (see the SUMMARY data set). • Verify that the entries required for the job are on the VIT. (All occurrences of VIT entries are listed at the bottom of the SUMMARY data set.) <ul style="list-style-type: none"> – Storage analysis requires the SMS VIT entries. – RU counting requires the PIU VIT entries. – VIT extraction origin and destination options work only on PIU VIT entries. • Check to see whether the VIT is a different level from the VIT analysis tool. For example, the DISP entry is now called DSP. Therefore, if you are extracting all occurrences of VIT option PSS with entry name e'DI*', no matches are found. • Do not specify a start or stop time. An event reported on the console can be off several seconds from the GTF time stamp. • Use the INTERVAL option to ensure that some output is seen before the job abends or is canceled.
Same output from a previous date	<p>If the job runs but the output data sets contain data from a previous job, check the DISP parameter. When DISP is NEW and the data set exists, the batch job runs anyway and then deletes the new data set. A message in the JCL log indicates whether this has happened.</p>
ABEND 80A	<p>If you are running the storage analysis function with MATCH and LENGTH options, try one or more of the following methods:</p> <ul style="list-style-type: none"> • Increase storage on the job. • Reduce storage pools to one storage type (GBLK, REQS, or VTAL). • Match only a few pools rather than all GBLK or all VTAL or all REQS pools. • Run the LENGTH option without the MATCH option, and save the output for future reference. The LENGTH option is independent of the MATCH option, so the same LENGTH output is shown, regardless of whether the MATCH option is specified. • Remove MATCH and LENGTH options. • Specify a start and stop time to limit the amount of data being processed.
Message CANNOT READ FILE WITH DD NAME TRACE received	<p>When processing multiple tapes using the VIT analysis tools, you receive the message CANNOT READ FILE WITH DD NAME TRACE and the return code is 10.</p> <p>If you are attempting to process multiple standard label (SL) tapes using the bypass label process (BLP), verify that the LABEL parameter on the TRACE DD statement is coded correctly. See Table 48 on page 649 to determine what document describes job control language (JCL).</p>

In addition to the actions suggested in Table 5 on page 45, try the following actions to help you diagnose the problem:

- Use the DEBUG option (add it as a keyword in the parameter data set), which produces large quantities of data showing what the tool is doing. Run this on a small portion of the trace to prevent the output from being too large to be useful.
- Run another tool, such as the IPCS GTFTRACE subcommand, or ACF/TAP to see whether they work on this trace data set and to compare output such as time stamps.
- Run a short job to see what is on the tape. A simple way to run a short job is to run VIT extraction with an expression that is never true. Use the NOFORMAT option to avoid the overhead of loading the format routine. You can use the following parameter data set as a short job:

```
Desc Tell what's in this trace data set by running without extracting
Desc any entries.
NOWRAP NOFORMAT
VITEXT e'zzzz'
```

This data set extracts all VIT entries with the name ZZZZ. (Presumably, no entries start with ZZZZ.) The SUMMARY data set shows whether the data set wrapped, which types of records are traced, which VIT entries and options are traced, and the first and last time stamps.

Documenting an APAR for VIT analysis tool problems

If an APAR is submitted, the following information is required:

- Input data
 - Format load module (AMDUSRFD)
 - JCL or REXX EXEC or CLIST
 - PARM data set
 - TRACE data set
 - VIT analysis tool load module (ISTRAFT1)
- Output data
 - JCL log
 - Data sets produced by the VIT analysis tool
 - DETAILS
 - LOG
 - OUTSTAN
 - SUMMARY
 - VITEXT

VTAM dump analysis tool problems

This information describes how to diagnose problems that might occur while running the VTAM dump analysis tools and includes the following topics:

- “Checklist for isolating the problem”
- “Documenting an APAR for dump analysis tool problems” on page 48

If you are having problems during installation of the tool, see z/OS Communications Server: New Function Summary for additional information.

Checklist for isolating the problem

1. Determine whether the error occurred as a result of an ISPF panel or a module. ISPF handles its own error conditions and displays them directly on the panel. If the error message appears in your IPCS output, it is probably issued from a formatted dump module.
2. If you submit your job using JCL, verify that no JCL errors are issued. If a bad return code is issued, determine whether it is a result of the JCL job or a formatted dump module.
3. If you receive the message Storage access failed for xxxxxxxx, browse the dump to determine whether the actual location exists in the dump. Storage

requests are usually in terms of the length of the control block. For example, if control block BB is X'20' bytes long, the storage service will be trying to retrieve X'20' bytes of data.

4. View the output to determine whether any error messages were issued during execution. Messages may indicate the cause of the termination.
If the VTAM formatted dump routine cannot access a field (either in the control block or in the chain of pointers to the control block), an abend will occur and a note of the condition is made on the dump output.
5. Check whether the required ISPF and IPCS maintenance has been applied as documented in the program directory. Check whether maintenance has been applied to IPCS, ISPF, or VTAM. If any of these are down-level, unpredictable results might occur.
6. If your ISPF prompt lists or PF keys are not working properly, see z/OS Communications Server: New Function Summary to ensure that everything is installed and concatenated properly.

Documenting an APAR for dump analysis tool problems

If an APAR is submitted for the problem, the following information is required:

- Dump used when error occurred
- IPCSPRNT output data set

IPCSPRNT is the output data set allocated by you to store all data generated during an IPCS session. See Table 48 on page 649 to determine what document describes IPCSPRNT.

- JCL if submitted through batch
- Maintenance levels of the following items:
 - FFST™
 - IPCS
 - ISPF
 - TSO/E REXX
 - VTAM

Knowing the level of IPCS, ISPF, and VTAM can help determine whether you are running back-level on these products.

Recommended documentation for VTAM problems

Symptoms are often related to a particular device, command, or update to the system. If you suspect this is so, tell the IBM Support Center of this relationship. The following information describes some possible relationships and the documentation you should have for each one.

APAR or PTF number

If the problem appears after you apply an authorized programming analysis report (APAR) fix, supply the APAR number. If the fix is a PTF (program temporary fix), supply the PTF number. The following table shows the format for APAR numbers and PTF numbers.

APAR	PTF
OAnnnnn OWnnnnn OYnnnnn	UAnnnnn UWnnnnn UYnnnnn

Device type

If the problem is associated with the use of a particular type of terminal or other hardware unit, supply that device type (such as 3278 Model 2). If the problem is associated with a particular type of communication link, supply appropriate link characteristics, such as SDLC, BSC, SNA, or non-SNA. Also, identify any recent microcode activity on the control units involved.

Operator command

If the problem is associated with a particular VTAM operator command, supply the full command name (such as VARY). Also, note any command operand (such as INACT) or a network node type (such as CDRM) that has been associated with the problem.

Terminal action

If the problem is associated with a particular terminal action, such as IBMECHO, USS LOGON, or pressing the CLEAR key, describe the action (or sequence of actions).

VTAM application program

If the problem is associated with a VTAM application program that is an IBM licensed program (such as CICS or TSO), supply the name of the licensed program.

Hardware error condition

Sometimes it is immediately apparent that a problem is related to a specific hardware error condition. The hardware error might have been detected and reported in several ways:

- By an operating system message
- By a VTAM or application program message
- By the system operator
- By a VTAM buffer filling up with information from one device
- Through LOGREC
- By a terminal user (an indicator of the error status is displayed in the operator information area, at the bottom of the terminal screen)

If a hardware error occurred, note the failure condition that accompanied it, such as UNIT CHECK or TIMEOUT.

If you think your problem is related to a hardware failure, use the following tools to collect information about the hardware failure:

- SDLC link level 2 (LL2) test. For more information about the LL2 test, see “Modifying SDLC link level 2 test” on page 179.
- NCP intensive mode error recording. For more information about intensive mode recording, see “Modifying NCP intensive mode recording” on page 179. See Table 48 on page 649 to determine what NCP document describes intensive mode recording.
- The NetView program, if you use it in your system.
- LOGREC (or similar operating system facilities).

- The VARY TCPIP,OSAENTA command if you think your problem is related to an OSA-Express2 or later failure; see z/OS Communications Server: IP Diagnosis Guide for more information about the OSAENTA command.

Note: For help with hardware problems, use the NetView program if you have it installed, or use the system console messages to identify the affected part of the network. If you need further assistance, contact your IBM branch office.

Coding change

A problem can occur after you make coding changes to the following things:

- VTAM network definitions
- Macro usage
- Start options
- User-coded exit routines
- Job control statements
- User applications

Supply information about the coding change. For example, if you change the PACING operand on an NCP LU definition statement, supply that information.

Use Table 6 to determine the type of documentation you need to either solve your problem or supply to the IBM Support Center.

Note: Documentation for the NetView program is included in Table 6.

Table 6. Recommended documentation for VTAM problems

Documentation	Description
Alias names	If your configuration is using SNA network interconnection and you are using alias names, keep a list of the alias names defined to each name translation program.
Application program log (if appropriate)	Some user-written operator application programs produce an application program log.
Exit routines	Keep a list of VTAM exit routines.
Link-edit map	If a VTAM load module is involved in a problem, an XREF map of the load module is needed to show the location of other VTAM modules within that load module. To get an XREF map, use the service aid LIST (AMBLIST) with the control statement LISTLOAD and the parameter OUTPUT=XREF. This produces a listing showing the module (CSECT) names and their location within the load module. See Table 48 on page 649 to determine what document describes how to use the LIST service aid.

Table 6. Recommended documentation for VTAM problems (continued)

Documentation	Description
Link Pack Area (LPA) map	Contains names and starting addresses of modules in SYS1.LPALIB. To get an LPA map, use the IBM service aid LIST (AMBLIST) with the control statement LISTLPA. See Table 48 on page 649 to determine what document describes the LIST service aid. When it is used with a link-edit map and a dump, an LPA map enables you to identify a module that is found at a specific address within the link pack area.
LOGREC	Contains records of various types of system failures, both hardware and software. For hardware failures, LOGREC entries contain sense and status information about the device causing the failure. For software failures, LOGREC entries contain information such as the program status word (PSW), the abend code, the failing module name (when possible), a symptom string, and the general registers at the time of failure. LOGREC entries are written each time VTAM produces a supervisor call (SVC) dump.
NetView hardcopy log (if using the NetView program)	Contains messages routed to the NetView program that are associated with an operator terminal.
NetView file (if using the NetView program)	<p>Contains session awareness data for all active sessions and session trace data for sessions with a resource for which a session monitor trace has been started.</p> <p>Session awareness data includes:</p> <ul style="list-style-type: none"> • Session type • Names of session partners • Session activation status • IDs of subarea physical units contained in the explicit route assigned to the session • Transmission group numbers • Addresses and network IDs of SSCPs that own links in the transmission groups <p>Session trace data includes:</p> <ul style="list-style-type: none"> • Session activation parameters • VTAM PIU data • NCP data
Network configuration	<p>List any application programs, new devices, or new levels of the operating system you have added to your network.</p> <p>Save the System Modification Program (SMP) configuration data set (CDS) for VTAM and TSO/VTAM components. See Table 48 on page 649 to determine what document describes SMP.</p>

Table 6. Recommended documentation for VTAM problems (continued)

Documentation	Description
Program Update Tape (PUT) and Program Temporary Fix (PTF)	<p>Supply a list of any PUTs and PTFs that have been applied to your system. Also, supply a list of changes that have been applied to the hardware, such as requests for engineering activity (REAs) and engineering changes (ECs).</p> <p>If you have identified a module as the source of the problem, supply the PTF eye-catcher if the module has one. (The PTF eye-catcher is the latest PTF number that has been applied to a module. It follows the module ID in a dump.)</p>
Routing data	<p>Keep a table of destination subareas, explicit route numbers, virtual route numbers, paths, and transmission groups as well as a table associating session types, Class of Service (CoS) names, and CoS tables.</p>
Symptom string	<p>Some VTAM routines provide a symptom string after a failure. After an abend, you will receive message IST931I, which contains the symptom string text. Refer to z/OS Communications Server: SNA Messages for a description of message IST931I.</p> <p>The symptom string is put in the system diagnostic work area (SDWA), which is printed by the Environmental Recording, Editing, and Printing (EREP) program as part of the LOGREC entries. See Table 48 on page 649 to determine what document contains more information on LOGREC.</p> <p>If a first failure support technology (FFST) probe produced the symptom string, EPW messages will appear in the console listing to describe the symptom string. See Table 48 on page 649 to determine what document contains more information on FFST messages.</p>
System-console hardcopy Log	<p>Shows all messages sent to or commands received from the operator. May help indicate when the system began to have problems. (VTAM problems may not be apparent at the time they occur.)</p> <p>If your installation has written its own version of a VTAM message, supply the original VTAM message when you report the problem.</p>
Tables	<p>Keep a list of the VTAM tables your installation has defined, such as USS and logmode.</p>

Table 6. Recommended documentation for VTAM problems (continued)

Documentation	Description
Version and Release number Component ID	<p>CSV1R9</p> <p>Component ID</p> <p>VTAM 5695-11701 Release 190</p> <p>At VTAM startup when VTAM initialization is completed, messages IST020I and IST1349I are issued with this information. Message IST020I displays the version and release number, and message IST1349I displays the component ID.</p> <p>In addition, information about the release level of each component is contained in an access-method-support vector list pointed to by the access method control block (ACB). See <i>z/OS Communications Server: SNA Programming</i> for more information about the ACB.</p>
VTAM definition library	<p>This is a set of definition statements for resources in the VTAM network, such as the application programs and network nodes. The VTAM definition library also contains the start options used to initialize VTAM, unless they were entered by the system operator. Include configuration lists and user installation exits with the definition library. Detailed information about the VTAM definition library is in <i>z/OS Communications Server: New Function Summary</i>.</p>

Methods for submitting documentation

You can send documentation to IBM using the following methods:

- File Transfer Protocol (FTP)
- Email
- Tape

Submitting documentation using FTP

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

Requirement: TRSMAIN is a prerequisite for PUTDOC.

To obtain PUTDOC and detailed instruction on its use, follow these steps in “Obtaining PUTDOC”:

Obtaining PUTDOC

These steps provide the minimum information that you need to obtain PUTDOC.

Procedure

Perform the following steps to obtain PUTDOC:

1. FTP to the website at <ftp://service.software.ibm.com>.

2. Log in using **anonymous** as the user ID and your email address as the password.
3. Change directories (cd) to the /s390/mvs/tools/putdoc/ directory, where you find three files: PUTDOC.BIN, PUTDOC.HTML and PUTDOC.SRC.
4. Read the PUTDOC.HTML file for detailed instructions.

Obtaining TRSMAIN

These steps provide the minimum information that you need to obtain TRSMAIN.

Procedure

Perform the following steps to obtain TRSMAIN and detailed instructions on its use:

1. FTP to the website at ftp://service.software.ibm.com.
2. Log in using **anonymous** as the user ID and your email address as the password.
3. Change directories (cd) to the /s390/mvs/tools/packlib/ directory, where you find two files: README.TXT and TRSMAIN.
4. Read the README file for detailed instructions.

Results

If you require any additional directions, call the IBM Support Center.

Using electronic transfer through email attachments

Smaller documents can be sent as attachments to an email message. This can include cutting and pasting user output or downloading the file to a workstation for inclusion. Displayable text can be downloaded using ASCII transfer; all others should be processed by the TRSMAIN utility (see "Obtaining TRSMAIN") and transferred in BINARY. Email systems usually 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).

Submitting documentation on tape

Whenever possible, submit documentation electronically. If, after talking to the IBM Support Center representative about a problem, you need to submit documentation to the VTAM service team and electronic submission is not possible, you can submit documentation on a tape. Documentation on tape can be handled most efficiently by the IBM Support Center if it conforms to the following guidelines.

Tapes that are submitted to the VTAM service team can be standard label (SL) or nonlabel (NL) cartridge (3480). Improved data recording capability (a feature on 3480, standard on 3490) (IDRC) can be used. Each tape should contain an external label to identify the tape and its contents in some way. If an APAR has been taken, put the APAR number on the label. Otherwise, put the PMR number on the label. If you use multiple tapes or multiple files on one tape, include a separate explanation itemizing the contents of each tape.

With each tape, include the output from the job used to create the tape. To verify that the tape was created correctly and that the job completed normally, the VTAM service team must have the output from the job that created the tape (not just the job control statements that were used).

Note: The information in this topic is from APAR OY17061. See that APAR for additional information.

To submit dumps, traces, and other information to the VTAM service team, take the following steps:

- For dumps

Do not format data in any way before or during the transfer of the dump to tape. Dumps can be transferred to tape using IPCS or IEBGENER. See Table 48 on page 649 to determine the document that describes how to use the IPCS and IEBGENER utilities.

Do not change the data control block (DCB) parameters of the dump data set. Define the DCB parameters as follows:

```
LRECL=4160, BLKSIZE=4160, RECFM=F
```

- For GTF traces

Move the GTF trace data from the trace data set (which is usually SYS1.TRACE) to tape using IEBGENER only. The DCB parameters for a GTF trace should be one of the following values:

```
LRECL=4092, BLKSIZE=4096, RECFM=VBA  
LRECL=4092, BLKSIZE=32760, RECFM=VBA
```

For both traces and dumps, do not reblock the data (that is, use a different BLKSIZE) when moving it to tape. Use only the DCB parameters shown in the preceding example.

Restriction: Using any other utility (IBM or non-IBM) to transfer dump or trace data to tape might result in a processing delay and result in the APAR being returned to you (closed “RET”) because the IBM service team is unable to process the tape.

- For other types of information

Other types of information (for example, VTAM definitions, NCP stage one input, and console logs) can be submitted on paper or tape. If you submit the data on tape, it should be written to tape using IEBGENER only. The DCB parameters used when writing this type of data to tape should be the same as the input data set (that is, the same DCB parameters as the source of the data).

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 VTAM 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.

If you have a complex problem, you might need to talk to several people when you report your problem to the IBM Support Center. Therefore, keep all the

information that you have gathered readily available. You might want to keep the items that are constantly required, such as the VTAM component ID, in a file for easy access.

Chapter 2. Collecting documentation for specific types of problems

After you have classified your problem as a specific type using information in Chapter 1, “Diagnosing VTAM problems: Where to begin,” on page 3, this topic shows you how to collect the additional information you need before contacting the IBM Support Center.

This topic includes the following information:

- “Common problem determination procedures” describes procedures for specific problem types.
- “Failing module” on page 100 tells what to do when you have isolated the problem to a specific module of the VTAM program. You might be sent to this information from within the procedure for the problem type you have chosen.
- “Symptom string structure” on page 102 describes the meaning of the fields found in a symptom string.
- “Reporting the problem to IBM” on page 103 describes how to report the problem to your local branch office or the IBM Support Center.

Common problem determination procedures

This information includes a description of the following procedures:

- “Abnormal end (abend)”
- “Wait” on page 61
 - “VTAM locks” on page 74
 - “Using the VARY INACT,FORCE command” on page 80
- “Loop” on page 81
- “Message problem” on page 87
- “Incorrect output” on page 90
- “Performance problem” on page 94
- “Storage problem procedure” on page 97
- “Documentation problem” on page 100

Abnormal end (abend)

If the problem is an abend, use the procedure in Figure 1 on page 58 to collect the following documentation:

- Chapter 5, “Using dumps,” on page 183
- LOGREC
- Symptom string
- Abend or system completion code
- Contents of the general registers (at the time of the abend)
- Module ID and PTF eye-catcher
- PSW (at the time of the abend)
- “Formatting and printing trace records” on page 323

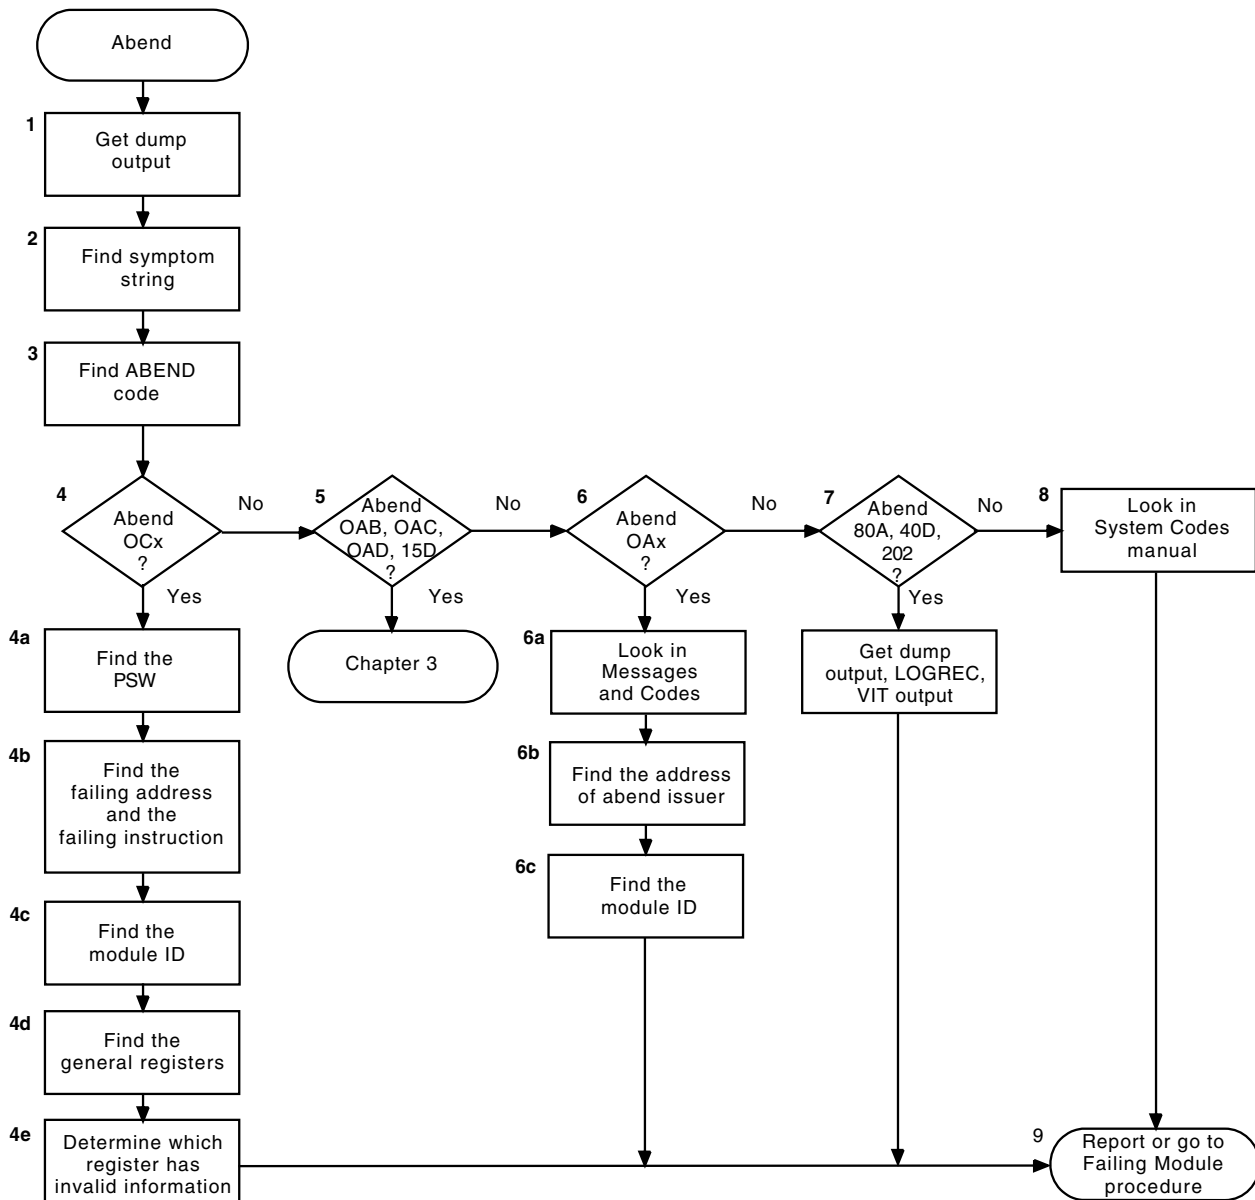


Figure 1. Overview of the abend procedure

The following procedure describe each step shown in Figure 1.

1. Get dump output. VTAM usually produces a dump for an abend. If no dump was taken, the dump files or spools might be full. Check for a message that an error occurred while VTAM was trying to produce the dump. If VTAM is not able to complete the dump, you have to re-create the abend or wait for it to occur again.

Note: To extract abend information from a VTAM dump, invoke the CLIST "ISTVABND" on page 242 or continue with the following steps.

2. Find the symptom string.

A symptom string is a structured database search argument. The symptom string gives information about what was happening at the time of the abend. Message IST931I, which contains the symptom string text, is issued when a symptom string is produced. A record is written to the LOGREC data set whenever VTAM takes a supervisor call (SVC) dump. For about 85% of all

abends, recovery routines produce a symptom string subset, which you can get by printing LOGREC. The symptom string subset, if it occurs, is located in control block SDWA in LOGREC.

The SDWA address should be listed in the beginning of the dump, in the dump abstract information. It is also printed out in LOGREC, labeled "Hex Dump of Record" at the end of each software entry.

The symptom string begins at X'194' in the SDWA. Field SDWAURAL gives the length of the symptom string, which can be up to 256 bytes.

3. Find the abend, system completion, or user completion code.

You can find the abend (or completion) code in the output of several different service aids. The system control block RTM2 work area (RTM2WA), the SYS1.LOGREC software record, and the task control block (field TCBCMPC) contain the completion code. The RTM2WA is pointed to by the TCB of the failing task (field TCBRTWA), and is listed after the abnormally ending TCB.

For more information on the abend codes issued by VTAM, see z/OS Communications Server: IP and SNA Codes.

4. Determine whether the abend code is OCx. If the completion code is of the form 0Cx (where x = the program interruption code from the PSW), continue with this step. If the abend code is not 0Cx, go to step 5 on page 60.

- a. Find the program status word (PSW) at the time of the abend.

The PSW is found in the LOGREC output, the SDWA, or the RTM2WA.

The location of the PSW in the dump output varies depending on the type of dump taken. For assistance in locating the PSW in dump output, see the diagnostic books for your operating system.

- b. Find the failing address or the failing instruction.

The PSW contains either the address of the next instruction to be executed at the time of the abend or the instruction that failed at the time of the abend, depending on the interruption code.

If the interruption code is X'10' or X'11', the PSW address points to the failing instruction. Otherwise, back up the PSW by the instruction length, and *that* is the failing instruction. Scan the dump output to find the address given in the PSW.

If you cannot find the address, the dump might not contain the relevant portion of main storage.

- c. Find the module ID for the module that contains the failing address.

VTAM identifies modules with the module name, Julian date, and PTF or APAR eye-catcher at or near the beginning of each module. This module identifier is in the form:

```
ISTxxxx yy.ddd nnnnnnn
```

where *xxxx* is the last five characters of the module name, *yy.ddd* is the Julian date the module was assembled, and *nnnnnnn* is the latest PTF or APAR fix (if any) that has been applied to this module.

Sometimes VTAM puts the module name of the failing module in LOGREC. If it is not there, you can find it in a dump. To find the module ID in a dump, start at the failing address and scan in descending address order along the right side of the listing. The module ID is printed in EBCDIC.

You can also scan the LPA map for the name of the load module and then go to the AMB list in the load module to find the CSECT that contains the failing address.

- d. Find the general registers.

The general registers in use at the time of the abend are found in the LOGREC output, the SDWA, or the RTM2WA.

Use the diagnostic books for your operating system to help find the registers.

- e. Determine which register has information that is not valid.

The failing instruction often uses a register with an address that is not valid in one of the general registers, or points to a location that is not valid (for example, low-address storage). Use *Principles of Operation* for your operating system, the program interruption code from the SDWA and the general registers used in the failing instruction, to determine (if possible) which register contains or points to incorrect data.

Note: When determining the validity of the register's contents, be careful to consider the address mode used by your operating system. Depending on the address mode being used, values used in 31-bit addressing might be interpreted differently than those used in 24-bit addressing.

Next go to step 9 on page 61.

5. Determine whether the abend code is 0AB, 0AC, 0AD, or 15D. These abend codes indicate a TSO/VTAM abend. For diagnosis information, see "TSO/VTAM abends" on page 109.

6. Determine whether the abend code is 0Ax. If the abend code is in the form 0Ax, continue with this step. If not, go to step 7.

- a. Find the abend code explanation in *z/OS Communications Server: IP and SNA Codes*.

An abend code of 0Ax indicates a problem within the VTAM network. The problem could have originated in VTAM, the NCP, an application program, or the hardware of some other network component. Look up the code in the information about Abend Codes in *z/OS Communications Server: IP and SNA Codes*. Most 0Ax abends place a return or reason code in register 15 at the time of failure. You can find the return code in register 15 by using the set of general registers from the LOGREC output, the SDWA, or the RMT2WA.

- b. Find the address of the module that issued the abend, using the PSW, which points to the next instruction after SVC 13.
- c. Find the module ID.

From the address determined in the previous step, scan in descending address order through the dump to find the module ID (see step 4c).

Go to step 9 on page 61.

7. Determine whether the abend code is 80A, 40D, or 202.

If the abend code is one of these, continue with this step. Otherwise, continue with step 8.

These abend codes indicate storage problems. Collect the following documentation:

- A dump of the VTAM address space
- A dump of the VTAM common storage area (CSA)
- LOGREC output
- VIT output at the time of the abend

After obtaining this documentation, go to "Reporting the problem to IBM" on page 103.

8. If the abend code is none of the above, see your operating system documentation.

To determine the publication that describes the abend codes for your operating system, see Table 48 on page 649.

Each code has an explanation of the documentation required and the problem determination steps to follow. For example, many abends occur during execution of SVC instructions. Parameter lists and register contents passed to SVC routines are in the diagnostic books for your operating system. These books might suggest that you obtain additional information such as a module name, a return code, a register containing information that is not valid, or the name of a system control block containing parameters that are not valid. After making a complete check of these sources, you are ready to report the problem.

9. Report or go to the failing module procedure. If you determined the module ID, go to "Failing module" on page 100. Otherwise, see "Reporting the problem to IBM" on page 103.

Wait

If the problem is a wait, use the procedure in Figure 2 on page 62 to collect the following documentation:

- "I/O trace" on page 340
- "Buffer contents trace output" on page 332
- Session trace data (if using the NetView program)
- Session awareness data (if using the NetView program)
- Dump of the VTAM primary address space including CSA
- List of:
 - Waiting process anchor blocks (PABs)
 - Waiting request elements (WREs) and associated event IDs (EIDs)
 - Waiting request parameter headers (RPHs)
- For problems associated with an application program:
 - "Formatting and printing trace records" on page 323
 - RPLs or FMCBs queued to the ACDEB
- For problems associated with the network:
 - Trace output
 - "Line trace" on page 354
 - "Generalized PIU trace" on page 353
 - "Transmission group trace" on page 359
 - "Scanner interface trace (3720, 3725, and 3745 only)" on page 359
 - Dump output
 - "Network control program (NCP) dump" on page 185
 - "Maintenance and operator subsystem dump (3720, 3725, and 3745 only)" on page 188
 - "Communication scanner processor dump (3720, 3725, and 3745 only)" on page 187
 - Reports from NetView, IMR, or EREP (if available)

Note: Use the documentation you have available to isolate or resolve the problem. If you have to re-create the problem, make sure the traces listed above are active.

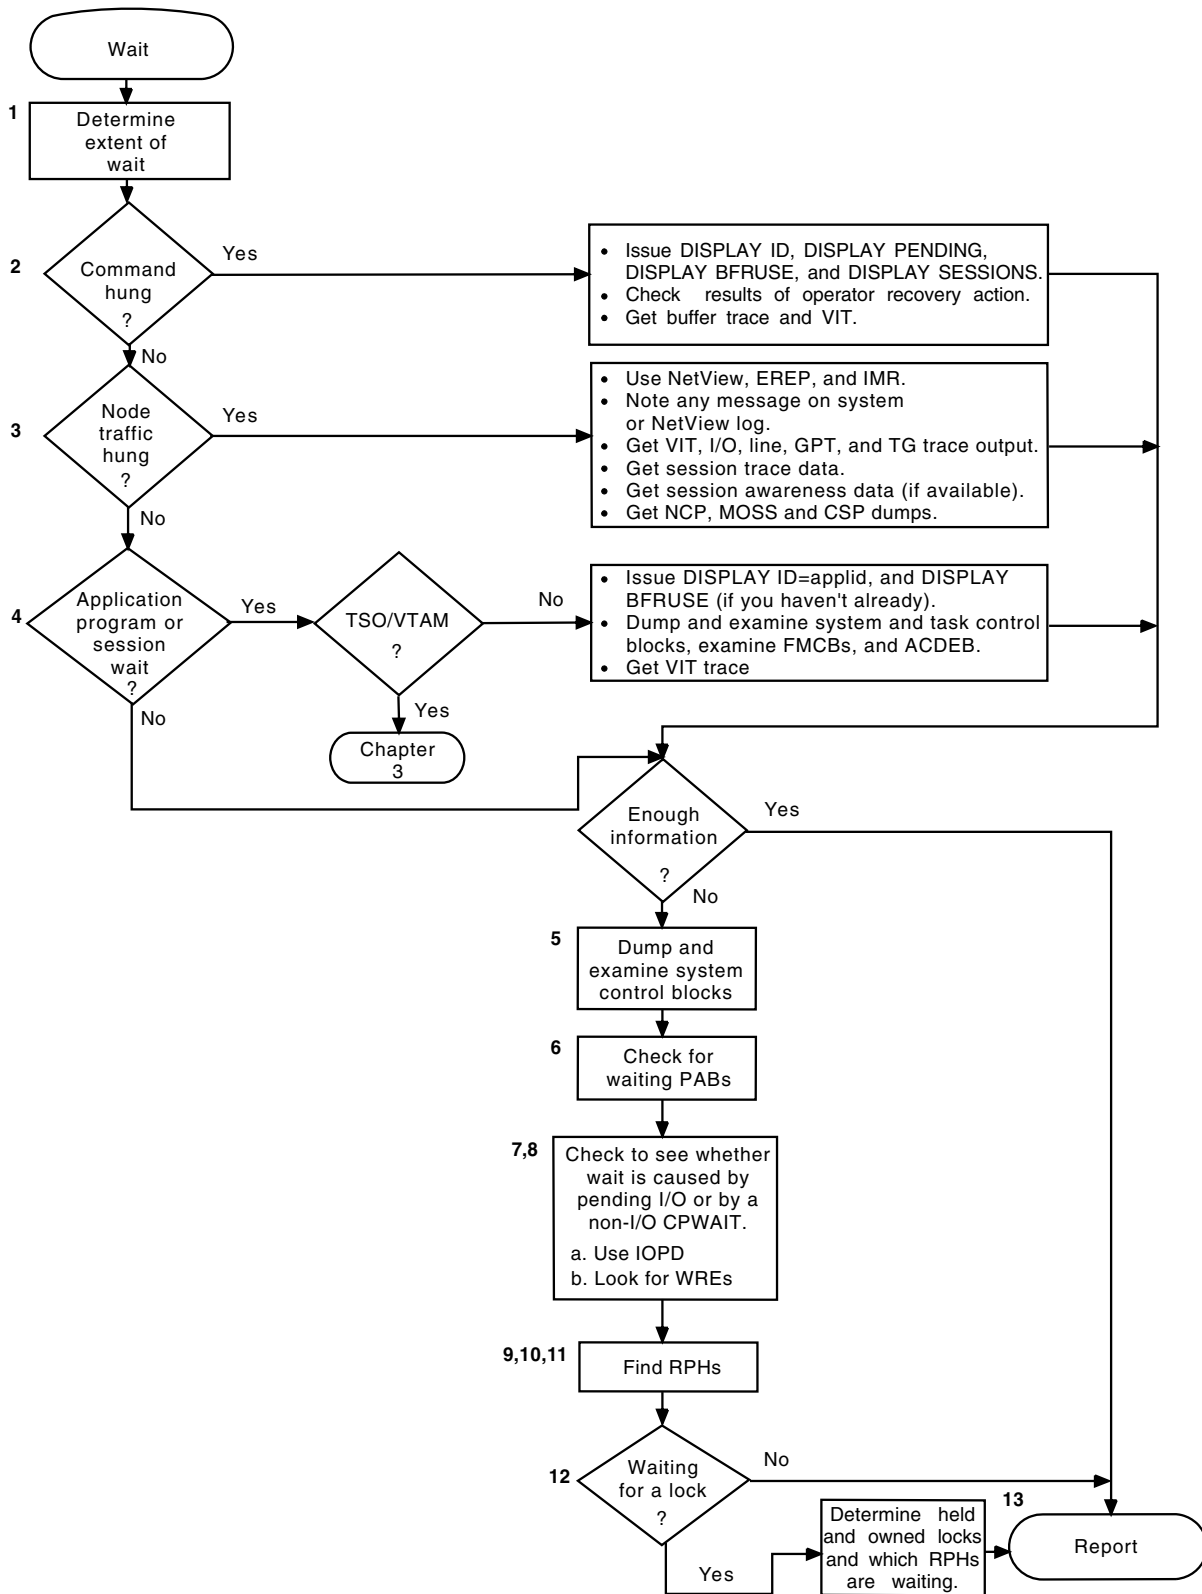


Figure 2. Overview of the wait procedure

The following procedure describes each step shown in Figure 2.

1. Determine the extent of the wait state.

Determine how extensive the wait state is in the operation of the VTAM network. Determine whether all VTAM processing stopped or only processing with respect to a single device, application, or something in between. Also determine what, if any, recovery action was taken at the time the wait was encountered by the operator or user. Some information about the activity that immediately preceded the wait might be available on the system log or in application program transaction logs.

2. Did a logon, logoff, or command fail to complete?

If so, continue with this step; otherwise, go to step 3.

- If the wait state was actually the failure of a VTAM procedure to complete, use the DISPLAY ID command to identify the status of VTAM resources at the time of the problem. Note any status codes that are abnormal.
- Use the VTAM DISPLAY PENDING, DISPLAY SESSIONS, or MODIFY IOPD commands to identify I/O requests for which VTAM is awaiting a response from a network node. Sometimes a network node appears in a pending state awaiting the completion of activity at a higher- or lower-level node (for example, PSUB1, PTRM2). The pending status on the other node is needed in such a case.
- Use the VTAM DISPLAY BFRUSE command to get information about VTAM buffer pools. Save the output for use later in this procedure.
- A VTAM operator might have attempted a recovery action (such as issuing a VARY INACT,FORCE command). “Using the VARY INACT,FORCE command” on page 80 shows how to determine whether this command completed. Check the node status to determine whether the recovery action reset the state of the node for which the original command was issued.
- If VTAM is waiting for an I/O response, look at the output of the VTAM buffer contents trace (assuming it is active when the problem occurs). If the trace shows that VTAM did send a request and is expecting a response, the problem is probably in another network node.
- You can get additional information about the status of a command from the VTAM internal trace (VIT). With the SSCP and PIU options, you can match requests and responses and determine any requests that are outstanding (that is, for which responses have not been received). The SMS option supplies information about resource usage, and the PSS option provides information about VTAM scheduling of the dispatching process. (See z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT for a description of the internal trace entries.)

At this point you might have enough documentation to report the problem to the Support Center. If so, go to “Reporting the problem to IBM” on page 103. Otherwise, go to step 5 on page 66.

3. Is network traffic stopped through a specific node?

If so, continue with this step. Otherwise, go to step 4 on page 64.

- Add the specific node type to your problem documentation. For example, the node could be a 3705, 3720, 3725, 3745, 3790, or a 3274. NetView and EREP facilities show whether errors have been recorded for the node in question. Session trace data (collected by the NetView program) shows whether the node is not responding to VTAM, or whether VTAM is discarding the responses. Consider using NCP intensive mode recording (IMR) for recurrent problems of this type.
- Note any messages on the system or NetView command facility log reporting ER-INOP outages or other failures. Use the VIT trace, or use the I/O trace with the EVERY operand, to trace the network flow up to the point of failure. NetView and LOGREC show the reason for the INOP.

- For NCP-related problems, use the line trace or generalized PIU trace if the affected node is in an adjacent subarea. Use the transmission group trace to record intermediate node flows up to the point where the problem occurred.
- If the problem might be in NCP software or communication controller hardware, obtain a dump of NCP storage. If the wait affects only part of the network, use the dynamic NCP dump facility. It allows the rest of the network to continue operating while the dump is taken. If the failure requires reactivating the NCP, use the MODIFY DUMP command. See “Network control program (NCP) dump” on page 185 for more information on NCP dumps.

If the NCP is hung or if the hung resource is attached to an NCP, see Table 48 on page 649 to determine what NCP diagnostic document describes troubleshooting the NCP.

- If the problem is in a channel-attached device or a channel-to-channel attachment, examine one of the following traces, if available, to determine the sequence of events preceding the wait. (If no trace output is available, you have to re-create the problem to get it.)

- VIT trace with the CIO option

- CCWTRACE

To determine what document describes I/O control blocks for your operating system, see Table 48 on page 649.

If enough information is available, go to “Reporting the problem to IBM” on page 103. Otherwise, go to step 5 on page 66.

4. Is it a session or application program wait?

If the wait state appears to be related to a particular VTAM application program, continue with this step. Otherwise, go to step 5 on page 66.

- Enter the DISPLAY ID command for the application program, using the EVERY or SCOPE=ALL operand. If there are any nodes with status ACT/U, reenter the DISPLAY command. If you are again informed that the status of a node is ACT/U, issue VARY INACT,FORCE for that node. If you still have a wait state, continue with the next step.
- If only one application program is waiting while others continue to communicate with VTAM, that application program probably contains an error. To determine what caused the problem, obtain a dump of the application program and the operating system supervisor at the time of the problem.
 - Make sure that the error is not an operating system error. (Use the diagnostic books for your operating system.)
 - If possible, use the dump to determine the reason the application program is waiting. If the application program is not waiting for VTAM, use the documentation for the application program to determine the reason for the wait. If the problem is in TSO/VTAM, see Chapter 3, “Collecting documentation for TSO/VTAM problems,” on page 105.
- If VTAM still seems to be the cause of the problem, you need output from the VIT to obtain a record of activity on the failing session. Because large amounts of data will wrap around in the internal trace table, you might want to specify MODE=EXT.

See z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT for more information on using the internal trace. You can also use the I/O or buffer contents traces to get information about all sessions with that application; specify ID=*application program name*.

- Using a dump of the problem, find the address of the VTAM ACDEB for the application program.

You can find an ACDEB associated with an application by using the VTAMMAP SES formatted dump tool. If VTAMMAP cannot be run, then find the ACDEB chain pointer in the ATCACDA field of the ATCVT.

- a. Use the ACDEB address to find it in the dump.

On the FMCB RECEIVE ANY queue, ACDRAFQH points to the first FMCB.

On the RPL RECEIVE ANY queue, ACDRARQ points to the first RPL.

Note:

- 1) If there are FMCBs (ACDRAFQH is not equal to 0), but no RPLs (ACDRARQ = 0), a problem has prevented the application program from issuing RECEIVES.
 - 2) If there are RPLs (ACDRARQ is not equal to 0), but no FMCBs (ACDRAFQH = 0), there might be a problem involving the continue any/continue specific (CA/CS) state of the session.
- b. Check for blocked PABs in the process scheduling table (PST). ACDTSKID points to the PST.

Look at the following PABs in the PST. To determine the offset locations for these PABs, see z/OS Communications Server: SNA Data Areas Volume 1.

PSTRQPAB

Request PAB

PSTRSPAB

Response PAB

PSTUEPAB

User exit PAB

See steps 6 on page 66 and 9 on page 73 for additional recommended actions.

- c. Get the LUCB address (field ACDLUCBA in the ACDEB).
- d. Get the address of a chain of FMCB extensions (field LUCFMCBA in the LUCB). Each FMCB extension represents one LU-LU session.
- e. Each FMCB extension contains a pointer (field TSPFMCBA) to the address of an associated FMCB. Find the FMCBs associated with hung sessions.

In those FMCBs, look for:

- The CA/CS indicator (in TSPPSFL1 and TSPPSFL2)
- The data queues (in TSPACCUM, TSPEWAIT, TSPNWAIT, TSPEDATA, TSPNDATA, TSPTSOP, and TSPTSIP)
- Session state flags (in TSPSESSR, TSPDTSR, TSPCRVSR, and TSPRQRSR)

- f. Determine whether there are any indications of unusual conditions. See z/OS Communications Server: SNA Data Areas Volume 1.
- g. Make a cross-reference listing of network addresses and node names to correlate the VIT PIU and I/O trace entries with VTAM session control blocks, such as the LUCB and FMCB.

See Table 48 on page 649 to determine what NCP document contains information on hung sessions.

If enough information is available, go to "Reporting the problem to IBM" on page 103. Otherwise, go to step 5.

5. Dump and examine the system data areas.

If you have not already done so, obtain a dump of the VTAM address space, CSA, LSQA, and SQA.

Find and analyze the task control blocks. Use the VTAMMAP PABSCAN dump tool to format the output. See "PABSCAN" on page 257 for information on using PABSCAN. See Table 48 on page 649 to determine what document contains more information on using dumps and finding and analyzing task control blocks.

6. Check for waiting PABs.

Note: You can use the VTAMMAP VTCVTPAB formatted dump tool as an alternative to step 6.

Look at the following PABs in the ATCVT. To determine the offset locations for these PABs, see z/OS Communications Server: SNA Data Areas Volume 1.

ATCCSPAB

Configuration services PAB

ATCVDPAB

VARY definition DYPAB

ATCPXPAB

Buffer pool expansion DYPAB

ATCPUPAB

Physical unit services DYPAB

ATCPUIOP

Physical unit services I/O DYPAB

ATCLUSRT

Logical unit services router DYPAB

ATCNTPAB

TSC no sessions DYPAB

ATCSSPAB

Session serialization PAB

ATCSOPAB

Session outage notification PAB

ATCCNSPB

CNS logon PAB

ATCTMPB

Message DYPAB

ATCTRMPB

Termination subtask DYPAB

Check the contents of the PABWEQP (or the PABVERYA for very extended PABs) and PABRPHA fields. The field PABWEQP in each PAB contains the address of a chain of work elements that have not yet been processed by VTAM. The field PABVERYA is defined at the same location as PABWEQA and contains a pointer to an array of WKE queues.

The array pointed to by the PABVERYA field contains the following information:

- A four-word header containing some control information about the very extended PAB.
- An array of work element queues in descending priority. For example, queue 1 is the first queue in the array, and it has the highest priority; queue 2 is the next queue in the array, and it has the next highest priority, and so on. Each queue has the following structure:
 - (Field PABVFRST) A pointer to the first WKE (head, or oldest) on this level queue
 - (Field PABVLAST) A pointer to the last WKE (tail, or youngest) on this level queue
 - (Field PABVSRVL) Service level
 - (Field PABVSRVC) Service count

The field PABRPHA in each PAB contains the address of an RPH that is either running or waiting.

Note: In some PABs, PABRPHA might contain the address of an RPH, even though the RPH is not running or waiting.

Note the contents of these fields in each of the PABs, and have this information available when you contact IBM.

Figure 3 on page 68 shows how to find each PAB. Figure 4 on page 68 shows the relative location of fields in a normal, extended, and slightly extended PAB. Figure 5 on page 69 shows the layout for a very extended PAB. The DYPAB begins X'10' bytes before the PAB.

Note: The PAB pointers shown in Figure 3 on page 68 are not contiguous in the ATCVT, but are shown that way for demonstration purposes only.

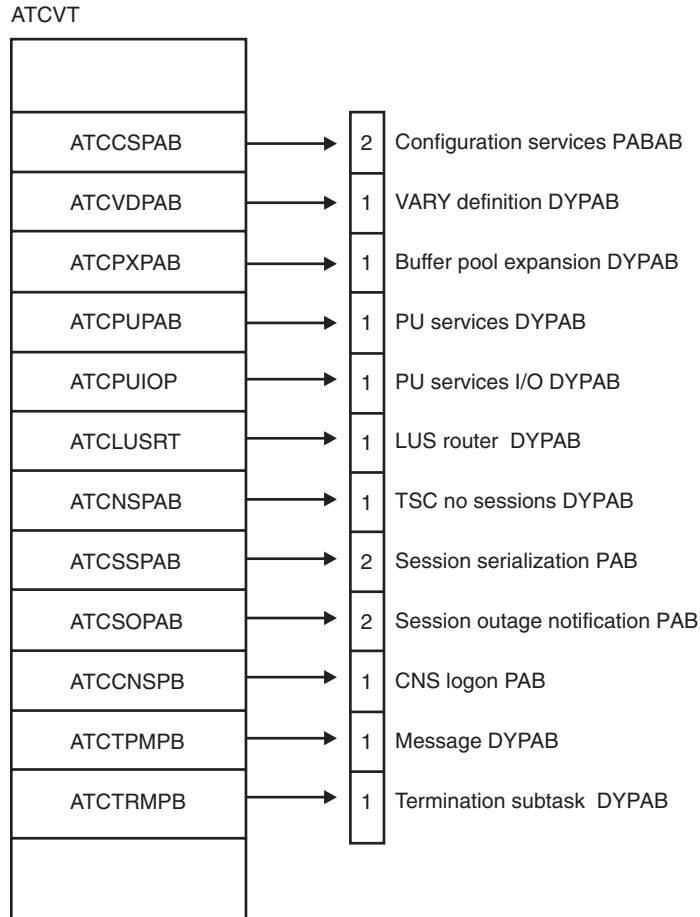


Figure 3. PAB locations

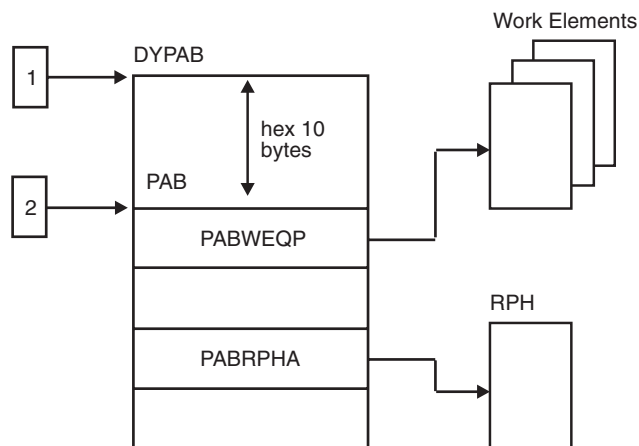


Figure 4. Normal PABs, extended PABs, and slightly extended PABs

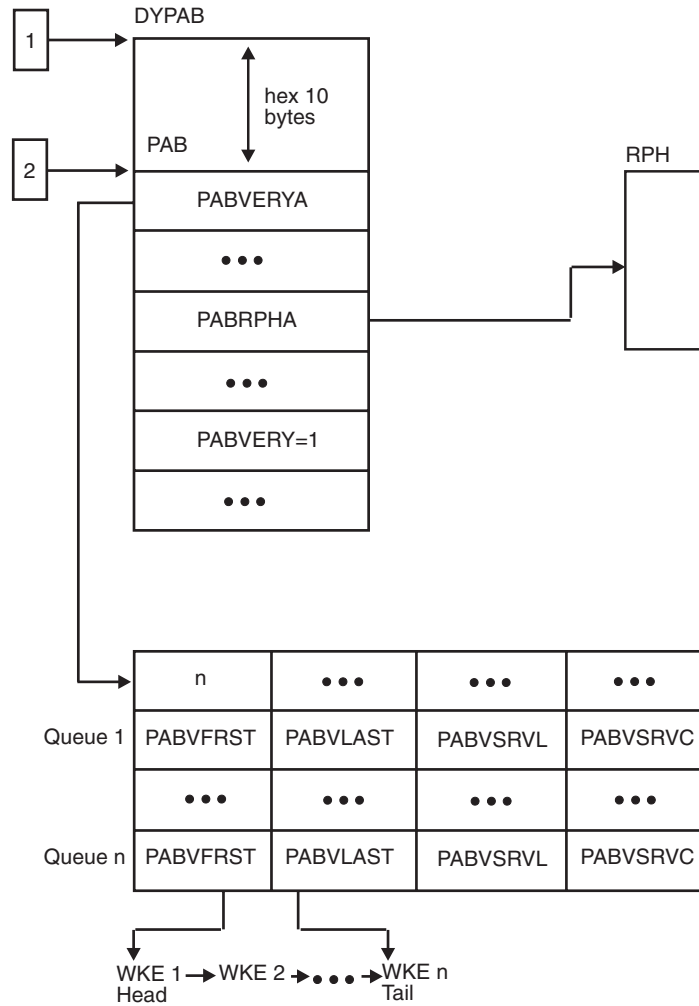


Figure 5. Very extended PAB

7. Is the wait caused by pending I/O?

Use the Input/Output Problem Determination (IOPD) facility to detect and report to the operator I/O operations that have been pending longer than a user-defined time limit.

When a VTAM process is waiting for a response, the process is represented by a waiting request element (WRE) queued to one or more LQABs within a single I/O LQAB group.

The WRE points to an event ID (EID), which indicates the reason for the wait. Look for the WREs and corresponding EIDs in a dump by using Figure 6 on page 72 and Figure 7 on page 73 and the following steps.

Note: You can use the VTAMMAP VTWRE formatted dump tool to count or help analyze WREs. See “VTWRE” on page 304 for information on using VTWRE.

- a. Find the address of the ATCVT at low-storage address X'408'.

If this low-address location is not available in a dump, use the pointer in the MVS control block CVT (CVTATCVT) to find the VTAM control block AVT. Location X'00' in the AVT points to the ATCVT.

The ATCVT is identified by release level at offset X'00' in the ATCVT. For z/OS Communications Server, the ATCVT is:

VE619(X'E5C5F6F1F9404040').

- b. Get the address of the I/O LQAB-group hash table from field ATCIOLQB. This hash table contains a number-of-entries field (LQHENTNM) followed by an array of table entries numbered starting with 0.

- c. Use the hash table to find the I/O LQAB groups for active subareas.

Each entry in the hash table is 4 bytes long and contains either 0, indicating an empty chain, or the address of the first LQAB group in a chain of I/O LQAB groups.

Within each I/O LQAB group, the LQGLINK field (offset X'10') contains the address of the next LQAB group in the chain. An LQGLINK value of 0 indicates the end of the chain.

- To find the I/O LQAB group for a specific subarea:
 - Calculate the hash table entry number, *N*, by dividing the subarea number by LQHENTNM and taking the remainder.
 - Search the chain for hash table entry *N* to find the LQAB group whose LQGSUBA field (offset X'0C') equals the subarea number.

Note: I/O LQAB groups are allocated only when needed. Therefore, you do not find an LQAB group for a subarea that has had no I/O traffic.

- To find all I/O LQAB groups, search the chain for each entry in the hash table.
- d. Find all the WREs chained off of a given I/O LQAB group.
 - Each I/O LQAB group contains several different LQABs. Use the global LQAB (LQGGLOBL) to analyze wait states, because its chain contains all of the group's WREs. (Chains off of the other LQABs in the group usually do not contain all of the group's WREs.) You can locate LQGGLOBL at the beginning of the LQAB group (offset 0).
 - The LQAB starts with the LQABFRST field, which contains either 0, indicating an empty chain, or the address of the first (oldest) WRE for this subarea.
 - Within each WRE, the WREGFWD field (offset 4) contains the address of the next WRE in the chain. The end of the chain is indicated by a WREGFWD value equal to the LQAB address minus 4.
 - e. Find the waiting event. Each WRE contains a WREIDCD field (offset X'32') that identifies the waiting event. The address and length of the waiting event ID are in the fields WREIDP (offset X'24') and WREIDL (offset X'30'), respectively.

For additional information, check the WREDTA field (offset X'2C'). In most cases, this field contains a CPCB operation code. If so, look in Appendix D, "Control point/control block (CPCB) operation codes," on page 623 to determine what function the operation code represents.

8. Is the wait caused by a non-I/O CPWAIT?

When a VTAM process has suspended itself using a CPWAIT and is waiting for a matching CPPOST or CPPURGE, the process is represented by a WRE queued to one or more LQABs within a single non-I/O LQAB group.

Analyze non-I/O CPWAITS using the steps described for pending I/O in step 7 on page 69, with the following exceptions:

- The IOPD facility does not detect and report these non-I/O events.

- No arrays or hash tables are used. Instead, each of the six LQAB groups is pointed to directly by its own address field in the ATCVT. These address fields are as follows:
 - ATCLUSMQ – logical unit services
 - ATCMCQAB – miscellaneous command
 - ATCPULQB – physical unit services
 - ATCNOSQ – network operator services
 - ATCSSLQB – SSCP session services 1
 - ATCSSMQB – SSCP session services 2
- WREs for non-I/O events do not contain a CPCB operation code value in the WREDTA field.

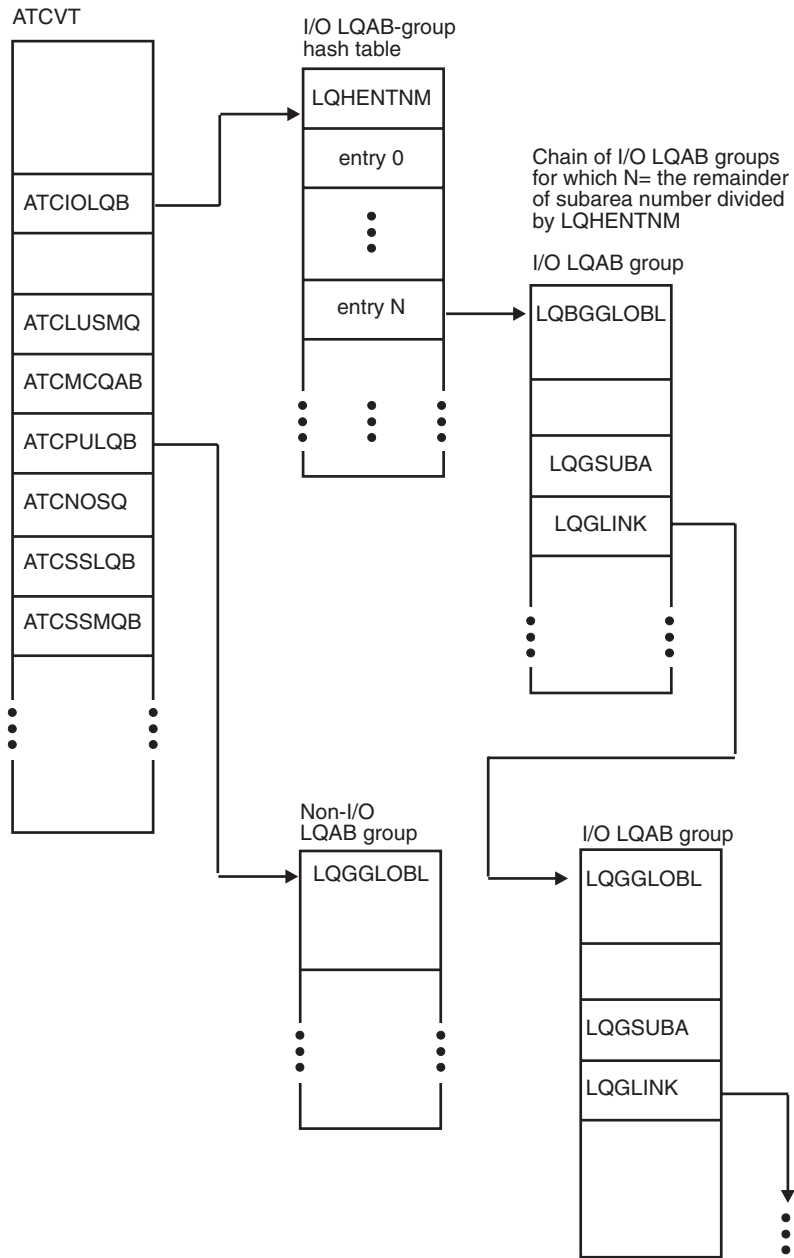


Figure 6. Finding LQAB groups

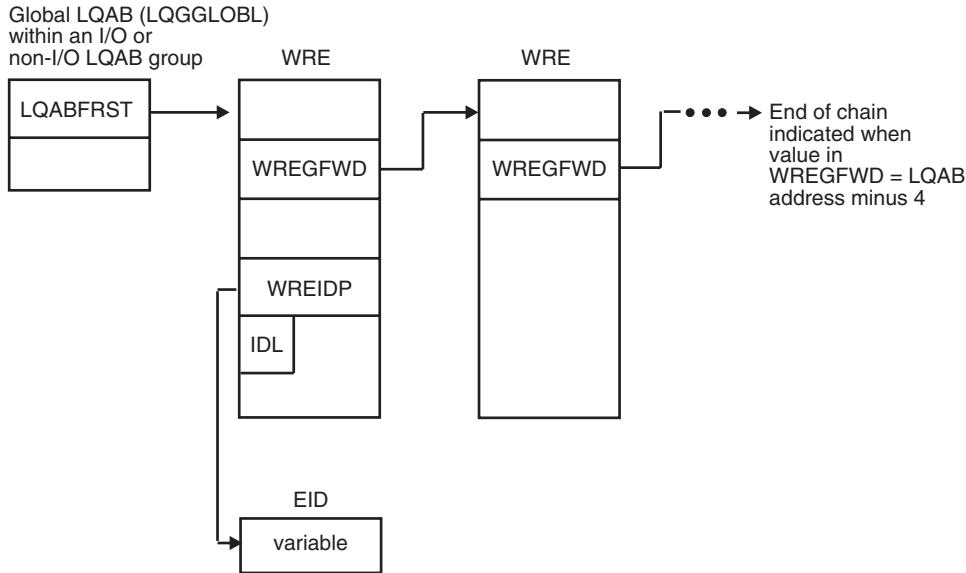


Figure 7. Finding waiting request elements for an LQAB group

9. Find waiting RPHs.

The following steps give instructions for examining two kinds of wait states: (1) a process waiting for a buffer, and (2) a process waiting for some other resource. Both kinds of waiting processes are represented by *request parameter header* (RPH) control blocks, but the RPH is found in different locations for each type of wait state.

- Step 10 explains how to find RPHs queued from a buffer pool control block. These RPHs show that the buffer pool cannot supply the required buffers, and as a result, the process is waiting. Note which buffer pool cannot supply the required buffers.
- Step 11 explains how to find RPHs that indicate a waiting process.

10. Find RPHs queued from buffer pool control blocks.

A buffer pool that has no available buffers can cause a wait state. There are many reasons for running out of buffers (for example, incorrect allocation in the VTAM start options, a VTAM programming problem, or an application programming problem). Use the DISPLAY BFRUSE output obtained in step 2 on page 63, if you were able to get it, to analyze buffer pool usage. Or use the VTAMMAP VTBUF and STORAGE formatted dump tools. See "VTBUF" on page 290 and "STORAGE" on page 282.

Also, follow the chain at offset X'04' into the RPH to obtain the addresses of other RPHs waiting for the same pool.

11. Find other waiting RPHs.

Waiting RPHs indicate a VTAM process that has not been completed. To locate the waiting RPHs, search the large pageable buffer pool (LPBUF) by hand or use the VTAMMAP VTRPH formatted dump tool. For more information, see "VTRPH" on page 299. Look at the formatted dump output.

Use the VTAMMAP VTBASIC formatted dump tool to analyze the request parameter headers (RPH) in the component recovery area (CRA). This function formats CRAs which contain RPHs. For more information, see "VTBASIC" on page 289.

12. Find RPHs waiting for locks.

- a. For each waiting RPH, look at the CRALxPTR fields. If any pointer (PTR) fields are nonzero, check the corresponding bit in CRALKACT. For example:
- If CRAL1PTR is nonzero, look at the last bit in CRALKACT.
 - If CRAL2PTR is nonzero, look at the next-to-last bit in CRALKACT.
 - If CRAL3PTR is nonzero, look at the third-from-last bit in CRALKACT.
- If the corresponding bit in CRALKACT is off (0), the RPH is waiting for this lock. If the bit is on (nonzero), the RPH is holding the lock and might be waiting for another lock. On your list of waiting RPHs, add the name of the lock being held or waited for. (See Table 7.)
- b. If you cannot find any locks waiting or being held using step 12a, scan the LPBUF buffer pool again, and list all allocated buffers that contain a nonzero value in field CRALKACT. These buffers indicate which RPHs own locks, if any, and which locks are held. A CRA can hold several locks. For example, a value of X'06' indicates two locks being held: the RDTLOCK (X'04') and the VOCLOCK (X'02'). (See Table 7.)
- For each allocated buffer with a nonzero CRALKACT field, look at the CRALxPTR fields. (The buffer might contain a resume address.) A nonzero pointer field contains a lockword address. Find the lockword. The first word of the lockword shows a queue of RPHs waiting for that lock. Add these RPHs to your documentation list.

13. Report the problem. Go to "Reporting the problem to IBM" on page 103.

VTAM locks

Table 7 includes a description of each VTAM lock, and Figure 8 on page 79 provides information on VTAM lock pointers.

Table 7. VTAM locks

Name	Lock ID	Lvl	Hex value	Control block	Field name	Quantity	Function
8SLOCK	1	3	04	MPNCB	MPN8SLK	One per multipath channel (MPC) line represented by an MPNCB	Serializes MPC outbound scheduling in a VTAM operating under MVS with multiple CPUs and System/390® or zSeries hardware. Ensures single remover for TPREMEL macros.
ADJLOCK	4	5	10	ADJSA	ATCADJLK	One per VTAM	Protects users of CIDCTL when adding or deleting an adjacent node.
AHHCLOCK	71	11	400	ISTTSEXT	TSEXT_LK	One per VTAM	Serializes access to the AHNCB queue in the TSEXT.
AHNCBLOK	31	5	10	AHNCB	AHNLOCK	One per active APPN host-to-host channel PU	Serializes AHNCB PU PAB with AHNCB PC PAB.
ASBREG	53	7	40	MNPS	MNPS_ALK	One per MNPS application	Serializes use of the pending registered CFS user's queue.
AULINLOK	63	8	80	AULIN	AULINLOK	One per VTAM	Serialize updates and references to list of Enterprise Extender lines.

Table 7. VTAM locks (continued)

Name	Lock ID	Lvl	Hex value	Control block	Field name	Quantity	Function
AUVTLOCK	68	7	40	ISTAUVT	AUVTLOCK	One per VTAM	Serializes access to two Enterprise Extender resources. One is a control block which represents a local IPADDR and the other is a control block which represents a resolved HostName.
BPBLOCK	38	3	04	BPB	BPBLOCK	One BPB per boundary function NCB	Protects BSB PCID and BSBSA tree for SNA/IP and rapid-transport protocol (RTP).
BSBLOCK	39	4	08	BSB	BSBLOCK	One per session using VTAM boundary	Protects updates and references of the BSB.
CIDLOCK	32	8	80	CIT	CITLOCK	One per session	Serializes changes to or deletion of FMCB.
CLKLOCK	69	9	100	ISTCLK	CLK_LOCK	One per VTAM	Serializes ISTRPCTM with HPRTIMER invokers.
CLWLOCK	70	9	100	ISTCLW	CLW_LOCK	One per VTAM	Serializes ISTAUCTM with IPTTIMER invokers.
CMMEMLOCK	48	10	200	CMDAT	CMMEMLPT	One per VTAM	Serializes access to list of large buffers allocated by CMIP.
CMPLOCK	46	5	10	CMPVT	CMP_LLNK	One per VTAM	Serializes access to list of active CMIP applications.
CMRPLOCK	45	5	10	CMDAT	CMRPLPTR	One per VTAM	Serializes access to data owned by CMIP replication and management information base (MIB) controller.
CONVLOCK	33	9	100	CONVT	CONVTLOCK	One per APPC conversation	Serializes deletions of RAB.
CRYPTOKLK	26	8	80	ATCVT	ATCRYKLW	One per VTAM	Serializes use of the session key token chain.
DEBX2LOK	27	6	20	DEBX	DEBX2_LK	One per ACB index table entry	Serializes queuing of an application API requests with the closing of an ACB.
DEBLOCK	6	5	10	ACDEB	ACDLOCK	One per OPEN application program	1. Protects FMCB queue off ACDEB. 2. Held by TSC and by OPEN or CLOSE.
DESCQLOK	62	7	40	INSTANCEDATA	DESCQ_LOCK	One per VTAM	Synchronizes removing of list descriptors from the list descriptor queue.
DWALOCK	15	8	80	DWA	DWALOCK	One per VTAM	Used by certain disabled TSC modules to serialize use of the disabled work area (DWA).

Table 7. VTAM locks (continued)

Name	Lock ID	Lvl	Hex value	Control block	Field name	Quantity	Function
FSEXTPLK	66	11	400	ISTFSEXT	FSEXTPLK	One per VTAM	Serializes queuing and dequeuing of the CFUSR block to PSTCFUSR queue. Serializes release of PST storage.
GENRSDEF	61	7	40	ISTGENRS	GENRS_LK	One per VTAM	Synchronizes queuing to the defer queue from RVM and the processing of the defer queue.
HITLOCK	41	3	04	HIT	HITLOCK	One per FID5 session address	Protects users of HPRCTL when assigning or deleting a FID5 address or when acquiring a BSB address through FID5 address lookup.
HNTELOCK5		7	40	HNTE	HNTELOCK	One per minor node (per host element address)	Serializes updates and references to control blocks based off the HNTE (RDTE, NCB, LUCB, FMCB).
HINTERBLK	23	8	80	HNT	HINTERBLK	One per minor node	Serializes APPC conversion data in the RAB.
HNTLOCK	7	6	20	HNT	ATCHNTLK	One	Protects updates and references to HNT during most CIDCTL functions.
HPRPSLOK	74	9	100	HPRPS	HPRPS_LOK	One per VTAM	Serialize access of the HPRPS control block.
HSQCHAIN	19	5	08	ATCVT	ATCHSQLK	One per VTAM	Serializes usage of the HSQH queues. One lock is used to protect all of the queues.
IAPTREE	40	2	02	SAACB	SAAIAPLK	One per VTAM	Serializes modifications and references to the IAP tree.
INNLOCK	17	9	100	ATCVT	ATCINNLK	One per VTAM	Ensures that PIUs that are going to a node that is in slowdown mode are sent in FIFO order.
IPNCBDIA	64	5	10	IPNCB	IPNCBDIA	One per VTAM	Serialize access to the list of dial-in lines for Enterprise Extender.
IUSAPLOK	57	11	10	ISTPST	IUSAPLOK	One per PST	Serializes APSINIT/APSTERM.
LKLNKSG	42	9	100	CMDAT	CMDLNKLNK	One per VTAM	Serializes access to certain control blocks of internal CMIP applications.
LMELLOCK	21	6	20	LME	LMELLOCK	One for every partner LU entry for every APPC application	Used to serialize access to partner LU information in the APPC logical unit mode (LM) Table.
LMHTLOCK22		5	10	LMHDR	LMHTLOCK	One per APPC application	Used to serialize access to the APPC logical unit mode (LM) Table.

Table 7. VTAM locks (continued)

Name	Lock ID	Lvl	Hex value	Control block	Field name	Quantity	Function
LSNLOCK	24	3	04	LSNCB	LSNLOCKW	One per PU connection to an IBM 3172 Interconnect Nways Controller	Serializes the LSNCB PU PAB with the LSNCB PC PAB.
LSVQLOCK	65	9	100	LSVT	LSVQLOCK	One per VTAM	Serialize access of LSNCBs pending deallocation queue.
LUTABLOK	43	8	80	LUTAB	LUTABLOK	One per slot in the LU/NCE hash table	Protects HPRCTL users when adding, deleting, updating, or finding LU entries in the LU/NCE table.
NCBQ	25	9	100	ATCVT	ATCLNLOK	One per VTAM	Serializes access to the queue of LSA NCBs anchored at ATCLNNCB.
NODATLOK72		9	100	NODAT	NODAT_LOK	One per VTAM	Serializes adding/deleting NODAT_EEDisplay control blocks on the NODAT_EEDisplayQ.
PDBUFLK	18	9	100	ATCVT	ATCBUFLK	One per VTAM	Allows the user to move in problem diagnosis trace data before the data is processed.
PSTIMERQ	59	7	40	CFSMNP	PSTQ_LOK	One per MNPS coupling facility structure represented by an MNPS structure object	Serializes use the outstanding PSTimer queue.
PSTLOCK	8	8	80	ATCVT	ATCPSTLK	One per VTAM	Serializes queuing and dequeuing of FMCB to PSTFMCB queue. Serializes release of PST storage.
QDCBLOCK	28	5	10	APNVT	APNQDCBL	One per VTAM	Serializes access to the queue of QDCBs attached to the APNVT.
QUEUE	16	9	100	PAB	PABLOCK	One per extended PAB	Serializes queuing and dequeuing of work elements to an extended PAB.
RDTLOCK	2	2	04	ATCVT	ATCRDTLK	One per VTAM	Protects users of CIDCTL (PAFIND). Obtained by PUNS when a network-addressable unit is to be added or deleted, or a use count decremented.
RMCBLOK	58	5	10	IUTRMCB	RMCSAPLK	One per VTAM	Serializes access to RM global IUSAP queue.
RMLCBLOK	56	7	40	RMLCB	RMLCBLOK	One per HPDT DLC	Serializes NCBCMPAB work queues.

Table 7. VTAM locks (continued)

Name	Lock ID	Lvl	Hex value	Control block	Field name	Quantity	Function
RPDCBLOK	55	7	40	ISTRPDCB	RPDCBLOK	One per RTP connection that a Performance Monitor (PMI) is monitoring	Serializes the adding and deleting from the unsolicited data queue for RTP path switch and RTP deactivation.
RPNPMILC	54	5	10	ISTRPNCB	RPNPMILK	One per RTP connection that a Performance Monitor (PMI) is monitoring	Serializes the collection of RTP data with the stop collection of data.
RTPHSQUE	49	5	10	RPNCB	RPN_HSLK	One per RPNCB	Serializes access to each rapid-transport protocol (RTP) NCB's (RPNCB) half-session queue.
RTPTBLOK	47	8	80	RTPTB	RTPTBLOK	One per slot in the rapid transport protocol (RTP) hash table	Protects HPRCTL users when adding, deleting, or finding RTPs in the RTP table.
RTPTBNLK	73	6	20	RTPTB	RTPTBNLK	One per VTAM	Serializes access to all ISTRTPNIs and their ISTFRTPs. Protects HPRCTL find with wildcard against HPRCTL add and delete.
SKTASGN	37	6	20	SAACB	SAA_ASGN	One per VTAM	Serializes assignment of sessions to the socket tasks.
SKTLOCK	36	8	80	SOTCB	SOT_LOCK	One per socket task (SOTCB)	Protects SOCCB chain off the SOTCB.
SLENTLOK	29	5	10	SLENT	SLE_LOCK	One per session list entry	Protects updates and references to the session list entry state indicators and to the sequential list of the TP work queue.
SSVCBLCK	67	11	400	ISTSSVCB	CFSSSVLK	1024 per Sysplexports structure	Serializes access to a list in the coupling facility Sysplexports structure.
TASKLOCK	35	7	40	SAACB	SAA_TASK	One per VTAM	Protects SOTCB chain off the SAACB.
TCEXTLOK	60	7	40	ISTTCEXT	TCEXTLOK	One per VTAM	Serializes access to TLNCB list.
TOKENCOL	52	7	40	TOKENCOL	COL_LOCK	One per collection object	Serializes access to the collection object.
TREELOCK	34	6	20	SAACB	SAA_TREE	One per VTAM	Protects the SOCCB tree.
VDLOCK	13	9	100	ATCVT	ATCVDLOK	One per VTAM	Serializes directed load processor.

Table 7. VTAM locks (continued)

Name	Lock ID	Lvl	Hex value	Control block	Field name	Quantity	Function
VOCLOCK	1	2	02	ATCVT	ATCVOCLK	One per VTAM	<ol style="list-style-type: none"> 1. Serializes OPEN/CLOSE with VARY. 2. Serializes VARY Activate, VARY Deactivate, and VARY ERP.
VRLOCK	20	3	04	VRBLK	VRBLOK	One per virtual route	Serializes usage of the VRBLK.
XCFCBLOK	50	9	100	XCFCB	XCFCBLOK	One (per VTAM)	Serializes access to the XCF NCB AVL tree.
XFNCBLOK	51	5	10	XFNCB	XFNCBLOK	One per other VTAM node in the sysplex	Serializes access to the XFNCB outbound data queues.
XHOTLOCK	44	9	100	CMPVT	XHOTLPTR	One per VTAM	Serializes calls to a nonreentrant module that allocates autodata for C PABs.

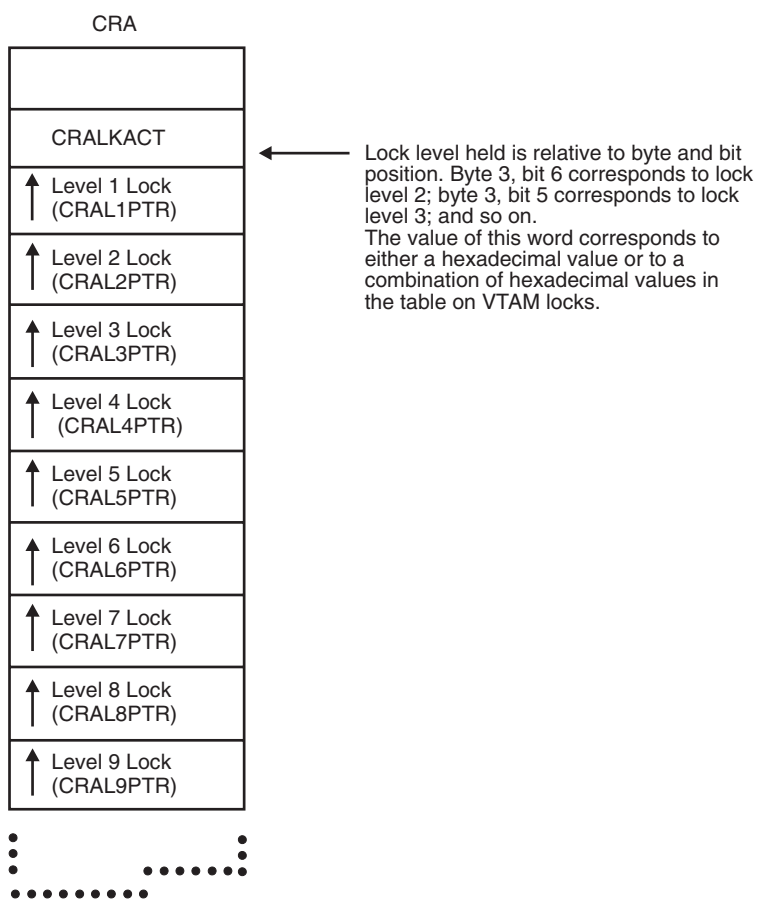


Figure 8. Pointers to VTAM locks

Using the VARY INACT,FORCE command

If the operator attempted a VARY INACT,FORCE command, check whether the command can be completed or whether there is a VTAM problem. Determine this using the following steps according to the resource specified on the VARY INACT command.

Note: Except for channel-attached SNA devices, if this command is issued to a resource with outstanding I/O, the command will not complete and VTAM must be recycled.

1. **Channel-attached physical unit or logical unit, SNA, or non-SNA:**
 - a. Display the resource status. If it is PHLIN, PHLAC, PDLUC, or PSUB1, the channel is hung or a required interrupt is missing.
 - b. If the status is PNFYx, go to step 11 on page 81.
2. **Link-attached SNA logical unit, switched logical unit:**
 - a. Display the resource status. If it is PNFYx, go to step 11 on page 81.
 - b. If it is anything else, there is a VTAM problem.
3. **Link-attached BSC 3270 logical unit:**
 - a. Display the resource status. You should see PDACL or PFDLU.
 - b. Issue VARY INACT,FORCE for the NCP or CA major node that defines the device.
 - c. Message IST105I indicates whether the deactivation succeeded and all lower-level nodes are inactive.
 - d. If the deactivation failed, display the status of all the resources in the NCP or CA major node.
 - e. If the status is PNFYx, go to step 11 on page 81.
 - f. If it is anything else, there could be a VTAM or NCP problem. Go to "Reporting the problem to IBM" on page 103.
4. **Link-attached SNA physical unit, switched physical unit:**
 - a. Display the resource status. You should see PDISC or PFDSC.
 - b. Issue VARY INACT,FORCE for either the physical unit to which the device is attached, or for the NCP or CA major node that defines the device.
 - c. Message IST105I indicates whether the deactivation succeeded and all lower-level nodes are inactive.
 - d. If the deactivation failed, display the status of all the resources attached to the NCP.
 - e. If the status is PSUBx, go to step 10 on page 81.
 - f. If the status is PNFYx, go to step 11 on page 81.
 - g. If it is anything else, there is a VTAM problem.
5. **Link-attached BSC 3270 physical unit:**
 - a. Display the resource status. You should see PDACP or PFDCP.
 - b. Issue VARY INACT,FORCE for the NCP.
 - c. Message IST105I indicates if the deactivation succeeded and all lower-level nodes are inactive.
 - d. If deactivation failed, display the status of all the resources in the NCP or CA major node.
 - e. If the status is PSUBx, go to step 10 on page 81.
 - f. If the status is PNFYx, go to step 11 on page 81.
 - g. If it is anything else, there is a VTAM problem.

6. **Local SNA or non-SNA major node, switched major node:**
 - a. Display the resource status. You should see PSUBx.
 - b. Issue VARY INACT,FORCE for any minor nodes that are not inactive. This should allow deactivation to be completed.
7. **Link:**
 - a. Display the resource status. You should see PDLNK.
 - b. Issue VARY INACT,FORCE for the NCP to which the link is attached. This should allow deactivation to be completed.
8. **Channel-attached NCP:**
 - a. Display the resource status. You should see PDISC.
 - b. Press the RESET LOAD button on the communication controller. This should allow deactivation to be completed.
9. **Link-attached NCP:**
 - a. Display the resource status. You should see PSUBx.
 - b. Display the status of the lower-level nodes.
 - c. If the status is PNFYx, go to step 11.
 - d. If the status is anything else, there is a VTAM problem.
10. **PSUBx status:**
 - a. Display the status of the lower-level nodes to find any pending states.
 - b. Deactivate any active or pending nodes. This should allow deactivation to be completed.
11. **PNFYx status:**
 - a. **For application programs with an NSEXIT exit routine:**
If the VARY INACT,FORCE command is unable to complete, there is a VTAM problem. Otherwise, deactivation should complete.
 - b. **For application programs with only a LOSTERM exit routine:**
 - 1) If the application program has issued a CLSDST macroinstruction, deactivation should complete.
 - 2) If the application program has not issued a CLSDST macroinstruction for the logical unit, issue a second VARY INACT,FORCE for the logical unit in question. If that does not correct the problem, you might need to cancel the application program to allow the deactivation to complete. (Canceling the application program terminates all of the LU-LU sessions with the application program.)
Coding the LOSTERM parameter on the APPL definition statement allows you to recover this type of hung resource without having to cancel the application.
 - c. **For application programs with neither exit:**
Deactivation does not complete until the application program issues CLSDST, the application program closes its ACB, or the operator cancels the application program.

Return to step 2 on page "Wait" on page 61.

Loop

If the problem is a loop, use the procedure in Figure 9 on page 83 to collect the following documentation.

Note: If you are using TSO/VTAM, use this procedure. You do not need to go to Chapter 3, "Collecting documentation for TSO/VTAM problems," on page 105.

- System console log
- Messages associated with the loop (if any)
- Failing module ID
- Dump of the VTAM address space that is looping
- Error file output (LOGREC)
- For a problem associated with a specific device:
 - "Formatting and printing trace records" on page 323
 - "I/O trace" on page 340
 - Session trace data (if using the NetView program)
 - Session awareness data (if using the NetView program)
 - NetView report (if using the NetView program)
 - "Activating network traces" on page 308

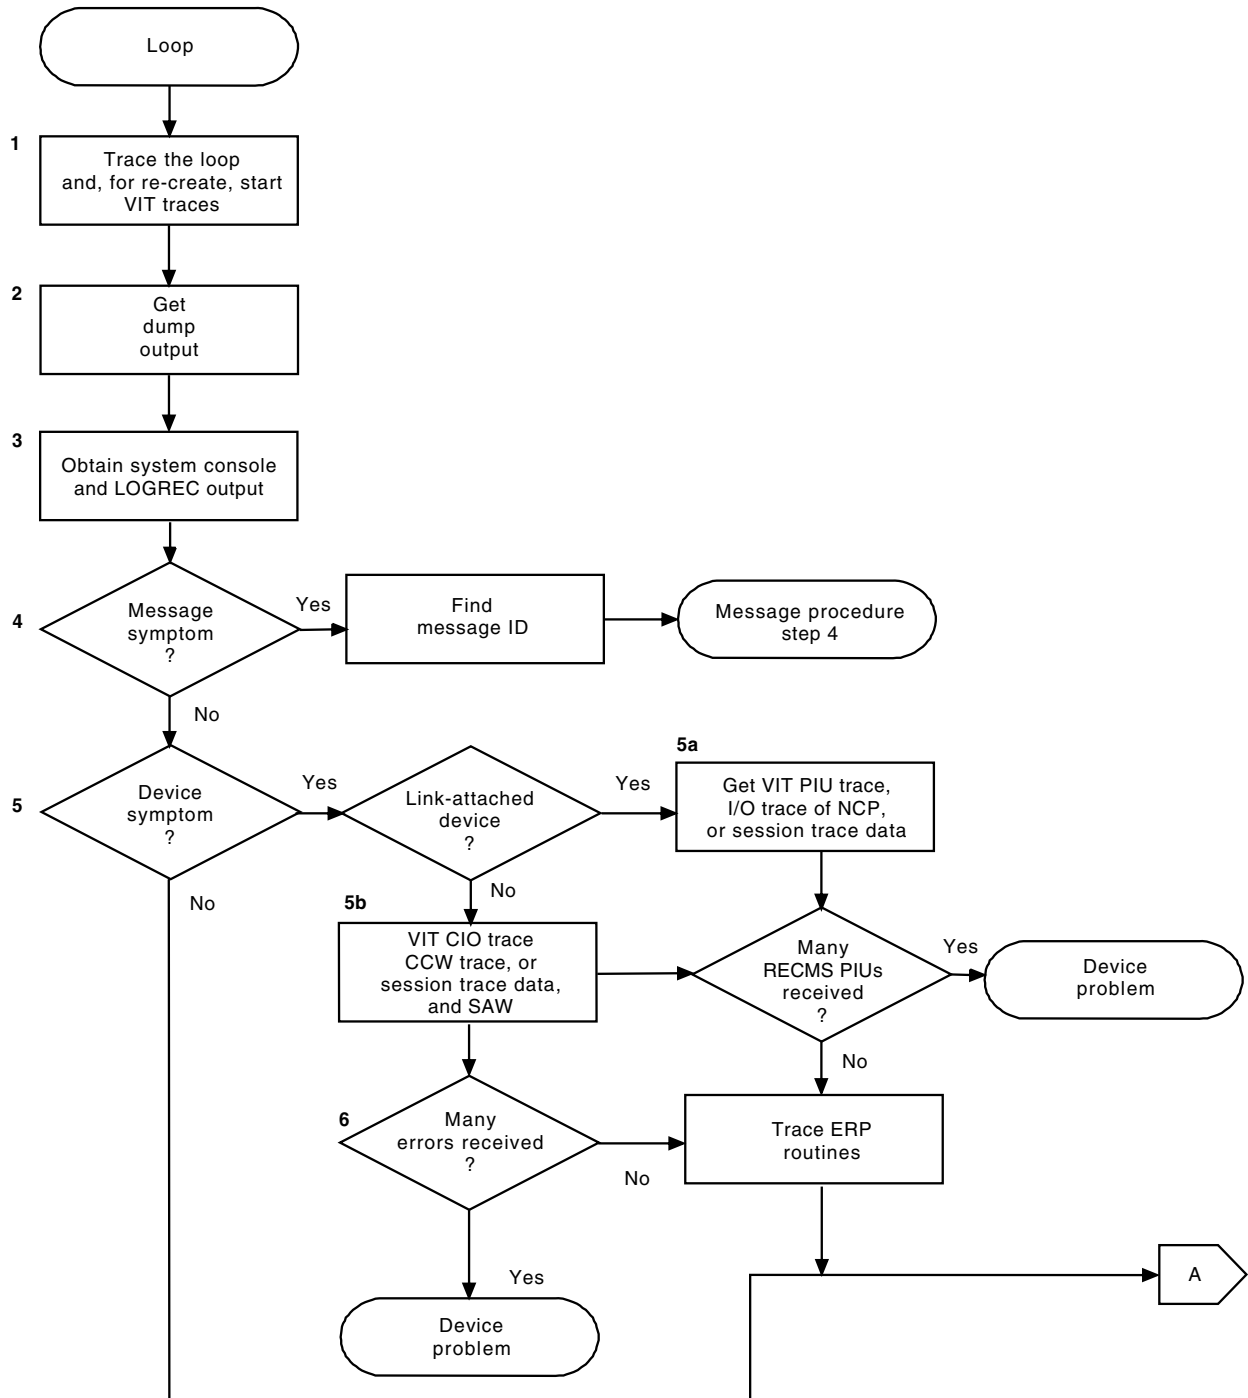


Figure 9. Overview of the loop procedure (part 1 of 2)

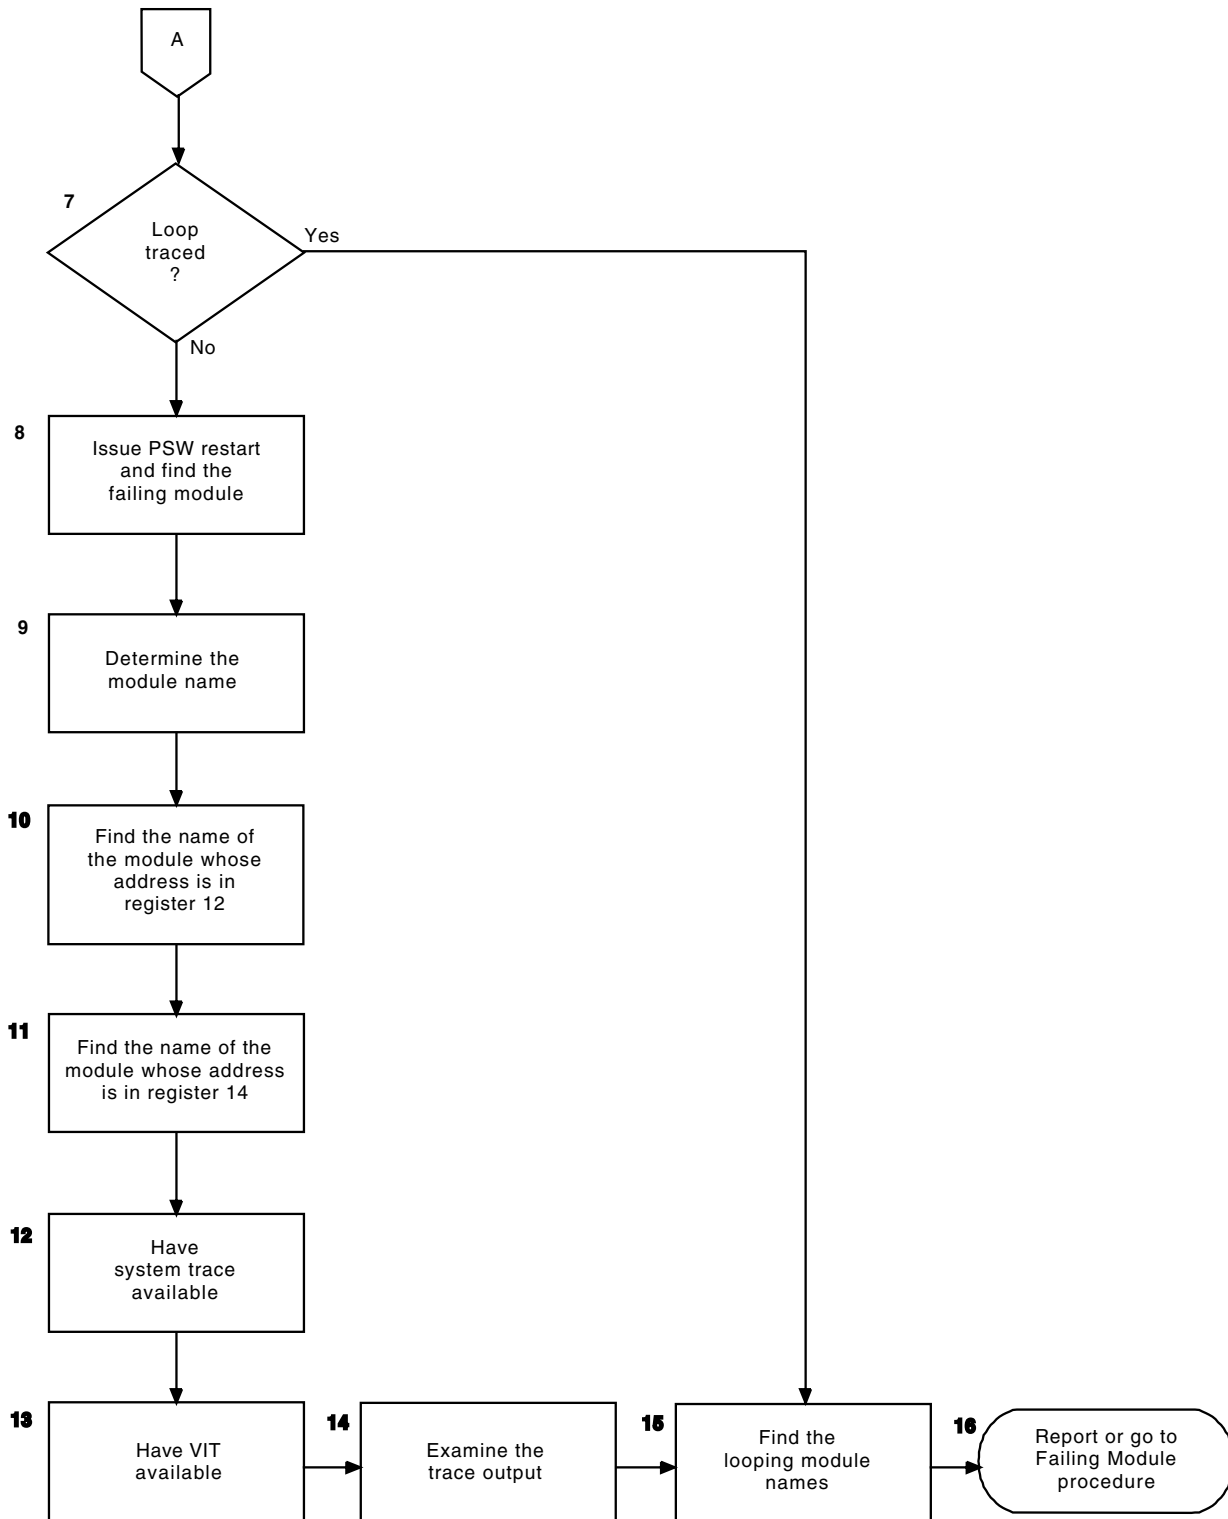


Figure 10. Overview of the loop procedure (part 2 of 2)

The following procedure describes each step shown in Figure 9 on page 83.

1. Trace the loop.

Loop problems might involve many modules or a single module. If possible, trace the looping instructions. Using the operator's reference for your host

processor, instruction-step through the looping addresses. Save these addresses for use in diagnosing the problem.

Take a dump and determine which module is looping by checking the PSW addresses in the CLKC entries for a repeating pattern.

If the VIT was running when the loop started, look for any exception conditions that might have led to the loop. If the internal trace was not running, you might have to re-create the problem to get the trace at the time of the loop. Set the internal trace to MODE=EXT to record the trace entries in an external file.

2. Get dump output.

To get a dump of VTAM, issue the DUMP command, or press the Program Restart key.

If the loop is disabled, the system console is not available for input, so take a stand-alone dump. (See "Stand-alone dump" on page 184.)

3. Get the system console log and LOGREC output.

The 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 Table 48 on page 649 to determine what document has information on LOGDATA.

4. Is a message involved?

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 numbers to your problem documentation and go to the message procedure, step 4 on page 89.

5. Is it a device error?

For any device error, first check the NetView report (if you have the NetView program) and then the LOGREC output.

Does the LOGREC output show repetitive entries for the same error on a particular device? If so, VTAM is receiving several different errors from that device.

- a. If the LOGREC error records are for a link or link station attached to a communication controller, get VIT PIU records and an I/O trace of the NCP. If you have the NetView program, get session trace data or session awareness data for the NCP. If the error records are for a link or device attached to a communication adapter, get VIT PIU records or a dynamic trace of the communication adapter.

If the trace shows continual arrival of RECMS PIUs, then the repetitive entries in LOGREC are caused by a device error.

Note: For information on counting PIUs see "Counting request/response units (RUs)" on page 377.

- b. For channel-attached devices, use one or more of the following traces for the device to determine whether VTAM is receiving many errors:
 - VTAM internal trace with CIO option
 - Session trace data (if using the NetView program)
 - Session awareness data (if using the NetView program)
 - CCWTRACE (if available)

6. Many errors received?

If VTAM is receiving many errors, the problem is probably in the device. Run a CIO VIT trace to trace execution of the VTAM ERP routines. Then continue with step 7.

7. Is the loop traced?

If you were able to instruction-step through the loop, go to step 15 on page 87; otherwise, continue with step 8.

8. Find the failing module.

Use the PSW to find the failing module.

- The PSW is found in LOGREC output, the SDWA, or the RTM2WA.

When you use PSW RESTART to terminate a looping task, a LOGREC entry is created with a completion code of X'071' for the task. An RTM2WA is also created for the task. Use the LOGREC record and the RTM work area to locate the failing module. See the diagnostic books listed in "Bibliography" for your operating system for help in locating the PSW in dump output.

Depending on the PSW bit 32, the last 3 bytes (24-bit mode) or 4 bytes (31-bit mode) of the PSW contain the address being executed at the time of the dump. Scan the dump output to find the address given in the PSW. See Table 48 on page 649 to determine which document contains more information on PSWs.

Note: Addresses might not always be in numeric order because the dump does not always generate output in sequential order.

If you cannot find the address, the dump might not contain the relevant portion of main storage. For example, the address might be in LPA storage. Have this portion of storage dumped, or use output from LPAMAP to identify the module, and proceed as above.

Note: The VTAMMAP VTFNDMOD formatted dump tool can be used to gather the module information described in steps 9, 10 and 11.

9. Find the module name that contains the failing address.

VTAM identifies modules with an EBCDIC module name and the Julian date (and, if appropriate, the latest PTF applied) at or near the beginning of most modules. This module identifier is usually in the form:

ISTxxxx yy.ddd [nnnnnnnn]

where xxxxx is the last five characters of the module name, yy.ddd is the Julian date the module was assembled, and nnnnnnnn is the latest PTF (if any) that has been applied to this module.

To find the module ID, start at the failing address and scan upward (in descending address order) along the right side of the dump listing. The module ID is printed in EBCDIC. Add the module name to your documentation list.

10. Find the module pointed to by register 12.

General register 12 (X'0C') is normally the base register for VTAM modules. In a VTAM loop, register 12 should point to the same module found in step 11. If not, add this module name to your documentation list.

11. Find the module pointed to by register 14.

General register 14 (X'0E') might point to a module that called the routine that is looping. Add this module name to your documentation list.

Add the module names from steps 9, 10, and 11 to your documentation list.

You can report the problem next, but you might need to continue with step 12 on page 87.

12. Get the system trace output.

The system trace might show many external and I/O interrupts. The PSW addresses in system trace entries will be part of the loop.

13. Get the VIT output.

The VIT is useful in determining the reason for a loop, such as a process being continually redispached for the same request. Get the VIT output. If you require VIT options in addition to the default options (API, CIO, MSG, NRM, PIU, PSS, SMS, and SSCP), start a VIT in addition to the default and specify MODE=EXT. If VTAM does not accept the command, it might be necessary to re-create the problem. For more information about using the VIT, see *z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT*.

14. Examine the trace entries.

By examining all of the trace entries, 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 VTAM system, or a control block chaining to itself. Check the output of the PSS option to see which VTAM routines are getting control. If you see a pattern of repetition in the trace entries, it does not necessarily mean that VTAM is looping. Some VTAM processes are timer-driven and repeat periodically.

Note:

- a. Get the trace information and examine the clock comparative entries for repeating PSW addresses. For short loops, the repeating PSWs show the extent of the loop.
- b. The absence of any apparent loop does not necessarily mean that VTAM is *not* looping. The loop might not contain a VTAM trace point.

If a module or modules are looping, get their addresses from the trace entries. Step 15 explains how to find the module name.

If you find a control block chained to itself, or if a queue of control blocks is in a cycle, try to identify the control block. Most control blocks have a 1-byte ID at offset X'00'. See the control block ID codes in Appendix E, "Storage and control block ID codes," on page 637 to identify the control block name.

15. Find the module names.

Note: You can also use the VTAMMAP VTFNDMOD formatted dump tool to find the module ID. See "VTFNDMOD" on page 294.

Use the addresses found in step 14 to find the module names involved in the loop.

To find the module ID, start at the failing address and scan upward (in descending address order) along the right side of the dump listing. The module ID is printed in EBCDIC. Add this module ID to your documentation list. Continue with step 16.

16. Report or go to the failing module procedure.

If you determined the module names, go to "Failing module" on page 100. Otherwise, you are ready to contact IBM. Go to "Reporting the problem to IBM" on page 103.

Message problem

If the problem is a message, use the procedure in Figure 11 on page 88 to collect the following documentation:

- Issuing module
- Message number
- System console log
- Chapter 5, “Using dumps,” on page 183
- “Formatting and printing trace records” on page 323

Note: If your installation changed the text of the message, the message ID might not be included, or might not match the ID of the message as it appears in z/OS Communications Server: SNA Messages. Therefore, it is recommended that you re-create the problem using the VTAM-supplied message text. Otherwise, determine what VTAM-supplied message text corresponds to the message text your installation is using.

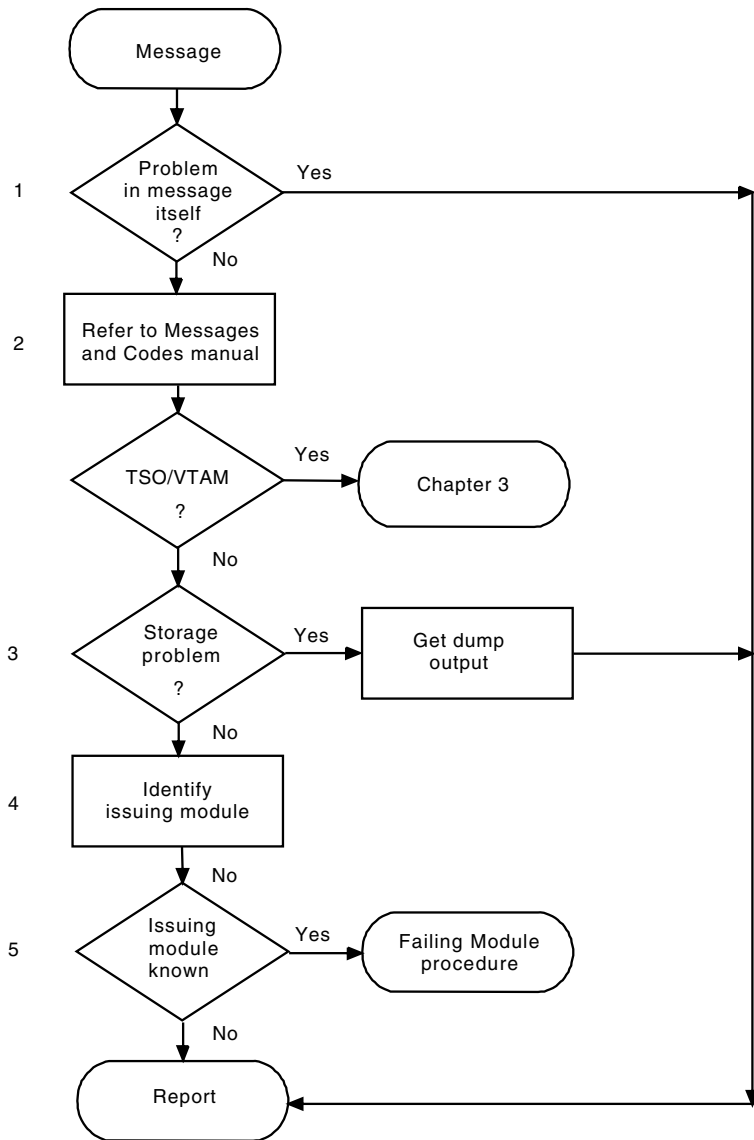


Figure 11. Overview of the message procedure

The following procedure describes each step shown in Figure 11.

1. Report if the problem is in the message itself.

If the content of the message is incorrect or the meaning of the message is not clear, go to "Reporting the problem to IBM" on page 103.

2. Follow the recommended action.

For all other messages, see z/OS Communications Server: SNA Messages for recommended operator and programmer actions. (See the list of VTAM books in "Bibliography" for the appropriate form number.) In addition:

- Identify the issuing component.
- If the message indicates a storage problem, go to step 3.
- If the message indicates a TSO/VTAM problem, see Chapter 3, "Collecting documentation for TSO/VTAM problems," on page 105.

The following list shows message prefixes and the components that issue those messages.

Prefix Issuing component

ELM	Logon Manager
IKT	TSO/VTAM
IST	VTAM
IUT	Connection Manager
IVT	CSM

Note:

- a. If the message starts with any other characters, it is issued from another network component or the operating system.
- b. Messages that begin with the prefix **ISTF** are issued by the VTAM dump analysis tools and the VTAM internal trace (VIT) analysis tool.

Help information for **ISTF** messages is available as a part of each tool by pressing F1. Therefore, these messages are not documented in z/OS Communications Server: SNA Messages.

See Chapter 6, "Using VTAM dump analysis tools," on page 191 and Chapter 8, "Using the VIT analysis tool," on page 365 for additional information about the dump and trace analysis tools.

3. Is there a storage problem?

If there is a storage problem, see "Storage problem procedure" on page 97 for additional information.

If there is not a storage problem, continue with step 4.

4. Identify the issuing module.

Try to identify the module issuing the message. If the MSGMOD start option is active or the MODIFY MSGMOD command is issued before the problem occurs, the message text contains the last five characters of the issuing module name. Add the message prefix to the module name, and add this name to your problem documentation. (To modify the module identifier in messages, see "Modifying message module identification" on page 178.)

The VTAM internal trace MSG entries contain the message number, the save area address, and the module ID (the 4th, 5th, 6th, 7th, and 8th characters of the module name). Use these to identify the issuing module. If the trace entry contains no module identifier, use the caller's address from the trace entry.

5. Report or go to the failing module procedure.

If you know the name of the issuing module, go to “Failing module” on page 100. If you are unable to determine the issuing module or resolve the problem, go to “Reporting the problem to IBM” on page 103.

Incorrect output

If the problem is *incorrect output*, use the procedure in Figure 12 on page 91 to collect the following documentation:

- Specific output that is incorrect
- Device type (if appropriate)
- “Buffer contents trace output” on page 332
- “Formatting and printing trace records” on page 323
- Session trace data (if using the NetView program)
- Session awareness data (if using the NetView program)
- “Network controller line trace (3710 only)” on page 358
- “Activating network traces” on page 308
- “TGET/TPUT trace for TSO/VTAM” on page 350
- Network problem:
 - “Line trace” on page 354
 - OSA-Express network traffic analyzer trace in z/OS Communications Server: SNA Network Implementation Guide
 - “Generalized PIU trace” on page 353
 - “Transmission group trace” on page 359
 - “Scanner interface trace (3720, 3725, and 3745 only)” on page 359

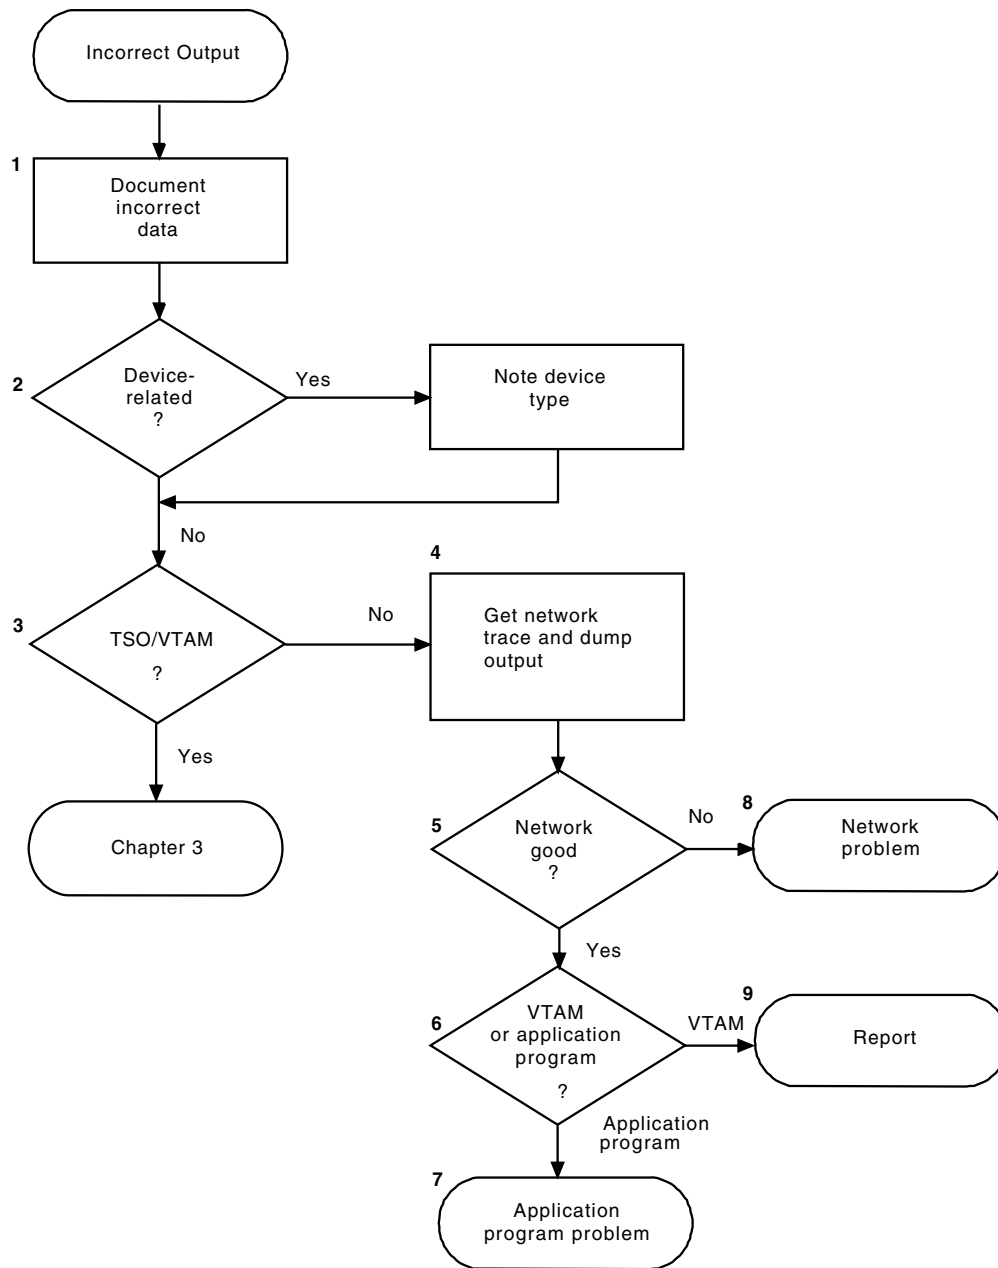


Figure 12. Overview of the incorrect output procedure

The following procedure describes each step shown in Figure 12.

1. Document the incorrect data.

Most incorrect output problems pertain to data contained in a PIU. This can be incorrectly formatted user data, routing information that is not valid, or other types of missing or incorrect data. These types of problems are generally difficult to diagnose, as they show up only at the user locations. From the following list, choose all the symptoms that apply to your problem and add them to your problem documentation:

- Cursor is in the wrong place or missing.
- Format of data is incorrect.
- Data is truncated.
- Data is incorrect.

- Data is missing.
- Problem is related to an application program macroinstruction.
- Screen is improperly formatted.
- Session is hung.
- Terminal is hung.

Note: The term *hung terminal* means that the user is prevented by the system from entering data.

2. Determine the device type.

If the problem is associated with a specific device type, add the device type (for example, 3277, 3278, or 3279 terminal) to your problem documentation.

3. Is it a TSO/VTAM user problem?

If the incorrect output problem involves TSO/VTAM, go to “Screen management problems” on page 112.

4. Get network trace and dump output.

Re-create the problem with the following service aids active:

- Start the VTAM buffer contents trace for the failing application program and terminal logical units.
- If you have an IBM 3710 Network Controller, start the network controller line trace. It traces information passing to and from a 3710.
- Start the VIT with MODE=EXT. Unless APPC is involved, do not specify the APPC option. Unless you suspect that a lock is not working, do not specify the LOCK option. Do not specify the MSG option. This shows the type of requests being processed between the application program and the user or terminal, and the control information for routing, pacing, and so on, in each PIU sent in the network.
- If the problem involves lines or devices attached to a communication controller, start the generalized PIU trace. This shows how far the PIU got within the NCP and what the PIU looked like (its control information) when it was sent to the line.
- If you have the NetView program, you can use the session trace data to determine the requests and responses received and sent by VTAM and the other network nodes.
- If it is available, you can use GTF CCWTRACE to trace the CCWs, I/O interruptions, and all CCW data for each Start I/O issued by the system. For more information about these traces, see the diagnostic books for your operating system.

Note: CCW trace will not capture data for a data device for the following devices:

- OSA-Express
- HiperSockets™

I/O trace must be used for these devices. CCW trace can be used for the control devices for the above devices.

- If it is available, you can use the VARY TCPIP,OSAENTA command to trace the packets sent to the network or received from the network by an OSA-Express2 or later adapter. The I/O trace captures the data as it is sent between VTAM and the OSA, but the OSAENTA trace captures the packets sent to and from the PCI bus on the adapter. For more information about the OSAENTA trace, see z/OS Communications Server: IP Diagnosis Guide.

- As soon after the problem occurs as possible, take a dump of the application program, VTAM, and TSO/TCAS. Stop all traces, and format the dump and trace output for online viewing.

The dump is used to reference storage addresses, such as control blocks and module entry points. The trace data shows at what point the data was modified, and what PABs the data was on as it was processed by VTAM. Take the dump during the re-create, when the traces are running. A dump taken earlier might not be accurate because the terminal device might have been deactivated and reactivated. This would allocate a different set of control blocks.

For more information on dumps and traces, see Chapter 5, "Using dumps," on page 183 and Chapter 7, "Using traces," on page 307. Operating system service aids are documented in operating system publications.

5. Examine the trace output.

Examine the individual trace entries to find the failure. If the problem concerns user data format, and the buffer contents trace or PIU trace does not show the incorrect data, use the output from the VIT trace with the SSCP option.

Use GTF CCWTRACE (if available) to see whether data is correct when it is sent to the NCP or logical unit.

Use the full buffer contents option for this trace. To use the full buffer contents option, specify AMOUNT=FULL on the buffer contents trace START option or on the MODIFY TRACE command. The VTAM internal trace records CC2, CI2, and CO2 contain the first 24 bytes of this data.

When output data is correct:

If the traces show that the data or the control information in the RH/TH as it leaves VTAM is correct, the problem is not in VTAM or the application program; go to step 9 on page 94. If the data going to the network is not valid, continue with step 7.

When input data is incorrect:

If the traces show that VTAM is receiving data that is not valid from a source external to VTAM, the problem is in the network; go to step 9 on page 94. If the data from the network is valid, the problem is in VTAM or an application program; continue with step 7.

6. Is it VTAM or an application program?

The problem has been narrowed down to VTAM or the application program. Examine each trace entry to determine whether the information from the application program was incorrect. If VTAM seems to be responsible, go to "Reporting the problem to IBM" on page 103; otherwise, continue with step 8.

7. Is more application program help needed?

For IBM application programs such as CICS or IMS[™], you can find additional diagnostic help in the IBM application program documentation. If you decide that the problem is with an IBM application program, contact the appropriate IBM representative for that product.

8. Is the problem with an external network device?

The problem has been narrowed down to the VTAM network, but not to VTAM itself. Try to identify the device or program responsible. You can use service aids, such as the NCP line trace, generalized PIU trace, or transmission group trace, to trace data flow between the NCP and terminal logical units. For information about how to use these traces, see "Traces provided by NCP" on page 353. For OSA devices, you can use CCW or I/O trace to trace data flow between VTAM and the OSA, and you can use the OSAENTA trace to trace

data flow between an OSA-Express2 or later adapter and the network. For a 3720, 3725, or 3745, use the scanner interface trace (SIT) to distinguish between NCP problems and line or terminal problems.

If you suspect the NCP, see Table 48 on page 649 to determine what document contains information on troubleshooting NCP problems.

Chapter 5, "Using dumps," on page 183 explains how to use system dumps, including the NCP dump. Contact the appropriate IBM representative for the device or program identified as the cause of the problem.

9. Report the problem.

Go to "Reporting the problem to IBM" on page 103.

Performance problem

If the problem is performance, use the procedure in Figure 13 on page 95 to collect the following documentation:

- System console log
- Error file output in LOGREC
- "Modifying tuning statistics" on page 181
- "SMS (buffer use) trace" on page 348
- "Network controller line trace (3710 only)" on page 358

Note: Performance problems do not generally indicate a VTAM problem.

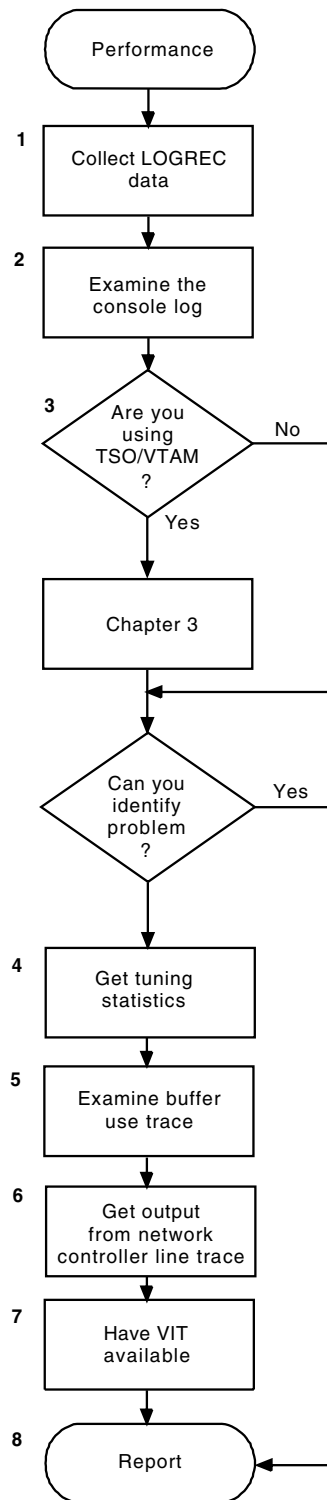


Figure 13. Overview of the performance procedure

The following procedure describes each step shown in Figure 13.

1. Get LOGREC output.

Performance problems are often caused by hardware errors. These hardware errors cause software error recovery processing to occur, which degrades system performance. For this reason, you should get the LOGREC output. LOGREC might show many hardware errors for a particular device or group of

devices. If the errors are limited to a single device, a hardware error is probably the cause. If the errors are displayed on many or all terminals of one type, software is more likely to be the problem, although hardware might still be at fault. If you suspect a particular device type, add it to your documentation list.

2. Examine the system console log.

The system console log might contain messages that help diagnose a problem. Add the message ID to your documentation list. The message prefix is IST, IUT, IVT, ELM, or IKT.

The system console log might also contain information about command problems. For example, operator commands might be taking too long to complete. Add the command name (for example, VARY ACT) to your documentation list.

3. For TSO/VTAM, see Chapter 3, "Collecting documentation for TSO/VTAM problems," on page 105.

If you are using TSO/VTAM, go to "Performance problems" on page 119. If you cannot resolve the problem with that procedure, return to this procedure.

4. Get tuning statistics.

If the performance problem is associated with traffic through a channel-attached host, a channel-attached communication controller, a channel-attached SNA physical unit, or multipath-channel-attached resources, it might be helpful to get tuning statistics for VTAM. (For more information about tuning statistics, see "Modifying tuning statistics" on page 181.)

5. Get output from the SMS (buffer use) trace.

You might have enough information to identify the problem. If so, go to "Reporting the problem to IBM" on page 103. If you do not, continue with this step.

a. Buffer pool expansion can cause performance problems. During VTAM initialization, error recovery, and VARY command processing, buffer usage is higher than normal. If buffer expansion is used, buffer pools should not expand except during such peak periods. Thus, what appears to be high buffer usage could be normal depending on the level of system activity.

Run the buffer use trace (TYPE=SMS). For information on how to start the trace and examine the output, see "SMS (buffer use) trace" on page 348. Coding the SNAPREQ start option causes trace entries to be written more often, providing a more comprehensive picture of buffer usage.

b. Using the time stamps in the system console and buffer use trace, correlate an excessive number of buffer pool expansions or large number of buffers used from a single pool with network activity recorded on the console. Constant high usage of a buffer pool might show that not enough buffers were allocated at VTAM initialization to properly support the level of network activity. Also look for a buffer pool that continually grows; buffers might not be released by some VTAM routines. Add the name of an active buffer pool (for example, LPBUF or IOBUF) to your documentation list.

6. Get output from the network controller line trace.

If an IBM 3710 Network Controller is installed, start the network controller line trace. This traces information passing over the lines to and from a 3710. [For more information about this trace, see "Network controller line trace (3710 only)" on page 358.] Print the trace output with TAP.

7. Get additional documentation.

If no solid indication of a problem is apparent, run the VIT with OPT=(PSS,API,SSCP,PIU) and MODE=EXT. This creates a history of VTAM activity. At the time of performance degradation, stop VIT and take a console dump of VTAM. (See your operating system manuals for information about how to take a dump.) Load the dump and trace output for future reference.

8. Report the problem.

Go to “Reporting the problem to IBM” on page 103.

Storage problem procedure

This procedure focuses on storage problems that occur in the common service area (CSA) or private storage area.

Procedure steps

The information in this topic is taken from the following VTAM storage diagnosis information APARs:

II06752

An Introduction/Overview

II04548

Documentation Requirements

II07563

Private Storage Problems

II07564

CSA Storage Problems

1. **Check for common CSA and private storage messages.**

Use the following messages to determine if the storage shortage is occurring in CSA or private storage. If the message is issued frequently or continuously, this indicates that a dump is needed to provide additional information.

Table 8 lists the messages that are associated with CSA storage problems.

Table 8. IST messages associated with CSA storage problems

Message number	Description
IST154I	Indicates that expansion failed for one of the fixed-length buffer pools in ECSA subpool 231. The error code displayed in the message provides additional information.
IST561I	Indicates that SLOWPT has occurred in one of the fixed-length buffer pools in ECSA subpool 231. <ul style="list-style-type: none"> • If this message occurs only occasionally, you might need to do some tuning in this area. • If the console is flooded with this message, you might have a CSA problem.

1. If you are running an LU 6.2 application, include the APPC VIT option in this list.

Table 8. IST messages associated with CSA storage problems (continued)

Message number	Description
IST562I	<p>Indicates that a storage request has failed because of one of the following reasons:</p> <ul style="list-style-type: none"> • CSALIMIT start option or MODIFY CSALIMIT is not specified and total CSA plus ECSA allocations have reached 90% of the total CSA plus ECSA defined in the system. • CSALIMIT start option or MODIFY CSALIMIT is specified and total CSA plus ECSA allocations have reached 75% of the total amount of CSA plus ECSA defined in the system. • CSALIMIT start option or MODIFY CSALIMIT is specified with the ,F command modifier and the total CSA plus ECSA VTAM usage has reached this value.
IST564I	<p>Indicates that a GETMAIN failed for the CSA subpool specified in the message. The specified subpool might be the sources of the CSA problem or the problem might be caused by another CSA subpool that is affecting the subpool displayed in the message. Further investigation into the contents of storage is required.</p>
IST1832I	<p>Indicates that the value coded on the CSALIMIT start option or a MODIFY CSALIMIT command is less than 25 megabytes, which may be too small. This might result in IST562I messages (only if the ,F command modifier is specified). Use the DISPLAY BFRUSE command output to monitor CSA usage and modify the CSALIMIT value if maximum CSA usage approaches the CSALIMIT value (this is necessary only if the ,F command modifier is specified).</p>
IST1833I	<p>This message will be issued only if the CSALIMIT start option or MODIFY CSALIMIT command is issued without the ,F modifier. Issue a DISPLAY BFRUSE,SUMMARY=* command and compare the MAXIMUM value in the first IST449I message to the value in IST1667I. If the first value is close to 75% of the second value, it means that VTAM is using a large proportion of system CSA storage. If so, an analysis of VTAM's storage use is indicated. If VTAM is not a large user of system storage but message IST1831I indicates close to only 25% of system CSA plus ECSA storage remains available for use, allocation of additional CSA or ECSA storage, or both, may be needed. A determination of which process is potentially using an inordinate amount of CSA or ECSA storage, or both, may also be needed.</p>

See the description of the message in z/OS Communications Server: SNA Messages for additional information.

Table 9 on page 99 lists the messages that are associated with private storage problems. See the description of the message in z/OS Communications Server: SNA Messages for additional information.

Table 9. IST messages associated with private storage problems

Message number	Description
IST563I	Indicates that the MAXPVT value has been reached. This value specifies how much private area subpool 229 storage VTAM can use within the address space of the application program displayed in the message. This indicates a problem with the application, not a VTAM problem.
IST565I	Indicates that a GETMAIN failed for the VTAM private area subpool displayed in the message. The specified subpool might be the source of the problem or the problem might be caused by another subpool that is affecting the subpool displayed in the message. Further investigation into the contents of storage will be required.
IST566I	This message is the same as IST563I except that MAXPVT was not specified on the APPL definition statement. This message does not indicate a VTAM problem.

2. Request a full dump.

If storage-related messages are issued frequently or continuously, dump VTAM common and private storage areas. The dump can help you determine the location of the storage problem.

- See “Formatting and printing dump output” on page 188 for information on the VTAM interactive problem control system (IPCS).
- Several storage-related dump analysis tools are available. See “STORAGE” on page 282, “VTAM” on page 287, and “VTBUF” on page 290 for descriptions of these tools.
- If external trace is active, see “Analyzing storage” on page 369 for information about analyzing storage using the VIT analysis tool. see information on internal and external trace recording in z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT for additional information.

Note:

- a. The best dump for diagnosing VTAM storage problems is a full dump.
If the dump is partial, examine the reason text of MVS message IEA911E to correct the problem so that you can obtain a full dump. The most common reason for a partial dump is that the dump data set is not large enough. In this situation, calculate the DASD space requirements and reallocate the dump data set.
For a complete description of the required documentation for storage problems, see information APAR II04548.
- b. Although VTAM detects storage shortages in the common storage area, VTAM might not be causing the shortage because this area is shared by all address spaces.

3. Use IBMLink to find additional problem determination information.

If you have access to IBMLink, take the following actions:

- Review the appropriate VTAM storage diagnosis information APARs. See “Procedure steps” on page 97 for a list of these APARs.
- Use your error messages and dump to determine key words for searching IBMLink for additional information and known problems.

If you do not have access to IBMLink and need additional assistance, go to step 4.

4. If you need additional assistance, contact the IBM support center at 1-800-IBM-SERV.

Documentation problem

Note: Before using this procedure, be sure that documentation is the problem. A VTAM problem might cause the documentation to appear wrong.

If the problem is documentation, use the procedure in Figure 14 to collect the following documentation:

- Incorrect information
- Form number of document

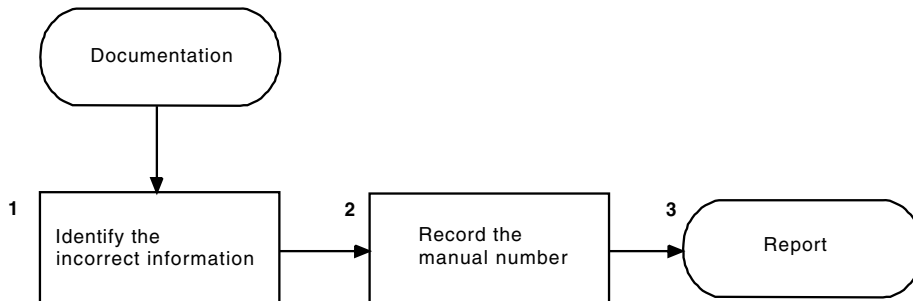


Figure 14. Overview of the documentation procedure

The following procedure describes each step shown in Figure 14.

1. Identify the incorrect information.

Add to your documentation list the name of the macro, operand, or procedure that is incorrectly defined or explained in the documentation (for example, *line trace*).

2. Record the form number.

Add the form number of the VTAM document to your documentation list in the form *ccnnnnnnrrr* (omit the dashes in the number; *rr* is the revision level). For example, report the form number of this document as GC31-6850.

3. Report the problem.

See “Reporting the problem to IBM” on page 103.

Note: Report a documentation problem only when it causes a VTAM problem. For suggestions, comments, or questions about z/OS Communications Server books, use the Reader's Comment Form at the back of the document.

Failing module

Use this procedure if you have identified a failing VTAM module in one of the other procedures (abnormal end, message, or loop). Figure 15 on page 101 shows an overview of the failing module procedure.

Use this procedure to get the following documentation:

- Module ID and PTF eye-catcher
- Caller of module
- “Formatting and printing trace records” on page 323

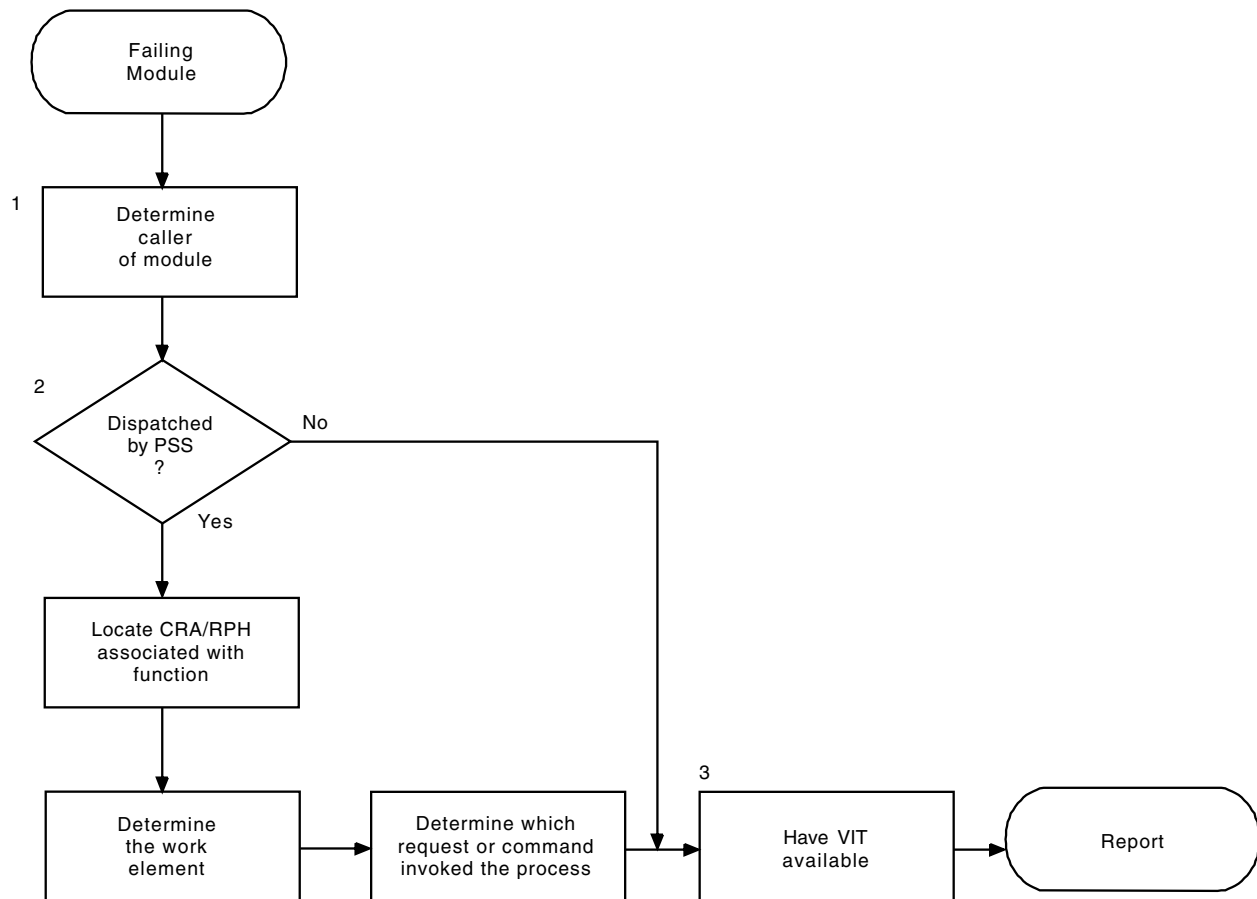


Figure 15. Overview of the failing module procedure

The following procedure describes each step shown in Figure 15.

1. Determine the caller of the service routine.

The failure might have occurred in a VTAM service routine used for many purposes. Determine the caller of the service routine. Use the save area conventions in “Using save-area module linkage conventions—Subarea” on page 405, or if you know the save area address, use the ISTVSAVE CLIST.

2. Examine the work element structure.

If your problem type is an abnormal end or loop, and the module is dispatched under control of VTAM PSS, find the CRA/RPH associated with the process. (See step 10 on page 73 in the wait procedure.) The RPHWEA field (at offset X'1C') usually points to the work element associated with the process at the time it was dispatched. To identify the work element, see “Using save-area module linkage conventions—Subarea” on page 405.

From the work element, it might be possible to identify an SNA request/response type, an operator command, or an application program request that ultimately caused the process to receive control. Add this request or command name to your documentation list.

If the RPHWEA field does not point to the work element, continue with step 3.

3. Have VIT information available.

Use the VIT options PSS and SMS to get more information about how the failing module received control or where the relevant control blocks are found. To obtain the address of the work element and the module name of the process entry point, use the last dispatch (DSP) entry for the failing process.

Symptom string structure

A symptom string is included in the dump for an abend, the dump for a program check, and the dump for a first failure support technology (FFST) probe point.

Message IST931I is issued for abend messages, and messages beginning with EPW are issued for FFST. Both the IST and the EPW messages contain the symptom text string. Figure 16 shows an example of a symptom text string for an abend dump. For information about FFST symptom strings, see *z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT*.

```
AB/S00C4 PIDS/5695-11701 LVLS/301 LVLS/91.025 RIDS/ISTCFCWM
RIDS/ISTAPCES#R RIDS/ISTCF2#L FLDS/PSW ADRS/000006C4
VALU/HB0044770 FLDS/POWPSSQ ADRS/00000F1C REGS/0C6A2 REGS/0A018
VALU/HE0044770 PIDS/568508501 PTFS/00000000 PRCS/00000010
```

Figure 16. Example of a symptom text string in output

The meanings of the fields are given:

- AB** Abend interrupt code, such as 0C4.
- ADRS** Offset into the failing VTAM module.
- DEVS** Device type and model related to the problem or generic device class.
- FLDS** Fields, control blocks, and DSECTs labels.
- LVLS** VTAM version and release level, or Julian date when the failing module was compiled. If LVLS appears twice in the symptom string, then it shows each of these values separately.
- MS** Program or device message.
- OPCS** Operation codes.
- OVS** Storage or core that was overlaid.
- PCSS** Statements, commands, JCL.
- PIDS** VTAM component ID.
- PRCS** Return code, status code, condition, feedback.
- PTFS** VTAM service level.
- REGS** The first two hexadecimal digits show the register number, and the next three hexadecimal digits show the displacement. The displacement value is the difference between the value of the PSW Instruction Address and the content of the register. Each REGS field is shown only if the value is less than the PSW, and if the difference is less than 4K.

If the REGS field has a value of X'FFFFFF', then no register contents are less than the value of the PSW Instruction Address and within 4K of the PSW Instruction Address.

If the REGS field has a value of X'FF000', the failing CSECT used relative branching support, so no register contents are less than the value of the PSW Instruction Address and within 4K of the PSW Instruction Address. CSECTs that use relative branching support have no base registers, so this symptom string can be expected.

If the REGS field has a value of X'FE000', then the value of the PSW Instruction Address is less than decimal 512.

- RIDS** One of three kinds of modules:
- Recovery module, if followed by #R
 - Load module, if followed by #L
 - CSECT name of the failing VTAM module, if not followed by anything
- SIG** System or device issued operator warning signal.
- VALU** Field value or overlay length.
- WS** Wait state.

Reporting the problem to IBM

For non-VTAM problems, call your IBM branch office. For suspected VTAM problems, do either of the following steps:

- If you have access to IBMLink, search for known problems in this area. If no applicable matches are found, report the problem to IBM by using the electronic technical report (ETR) option on IBMLink.
- Contact the IBM Support Center at 1-800-IBM-SERV.

After asking for your account name and other customer identification, the service representative will ask for a brief description of the problem. Your documentation list should contain the answers to all questions related to the problem.

Chapter 3. Collecting documentation for TSO/VTAM problems

This chapter shows you what documentation to collect for each type of common problem with the TSO/VTAM program. Use this chapter with Chapter 2, “Collecting documentation for specific types of problems,” on page 57.

Note: Most traces discussed in this chapter are described in Chapter 7, “Using traces,” on page 307. The exceptions are SVC 93 and SVC 94 entries. See Table 48 on page 649 to determine what document describes the SVC 93 and SVC 94 entries. For VTAM and TSO/VTAM command syntax, see z/OS Communications Server: SNA Operation.

Initial TSO/VTAM problem analysis

To use this chapter, start below and follow the steps.

1. Are you receiving one or more of the following messages?
 - USS message 7 ‘LU-name UNABLE TO ESTABLISH SESSION — RU-name FAILED WITH SENSE sense’, or similar USS message
 - USS message 10 (the user-defined logon message)
 - Message IKT029I or IKT028I at the operator's console
 - Message IKJ608I at the operator's console. See Table 48 on page 649 to determine what document describes message IKJ608I.If so, go to “Logon problems” on page 106.
2. Have you encountered one of the following abends?
 - ABEND0AB
 - ABEND0AC
 - ABEND0AD
 - ABEND15DIf so, go to “TSO/VTAM abends” on page 109.
3. Are you having parameter initialization problems?
 - Message IKT013I or IKT014I at the operator's console.
 - Initialization parameters have not been used.If so, go to “Parameter initialization problems” on page 110.
4. Do you have a hung terminal?
 - The terminal does not respond to any keys you press.
 - You must enter data from the terminal before processing will continue (in a situation where output is expected).If so, go to “Hung terminal problems” on page 111.
5. Are you having one or more of the following screen management problems?
 - Data is in the wrong place on the screen.
 - Data stream errors occur (such as operation checks, commands are rejected, PROGxxx).
 - Function errors occur (such as incorrect full-screen processing or incorrect line prompting in input mode of TSO EDIT).
 - Data length is incorrect.
 - Data content is incorrect.

If so, go to “Screen management problems” on page 112.

6. Are you having one or more of the following screen size problems?
 - The terminal does not operate in the expected screen size after logon.
 - The screen is not always the expected size during a TSO session.

If so, go to “Screen size problems” on page 117.

7. Are you having one or more of the following performance problems?
 - Slow response time
 - An increase in the number of detected waits
 - An increase in the number of swap-outs

If so, go to “Performance problems” on page 119.

If your problem is not listed in the steps above, it is probably not a TSO/VTAM problem. Go back to Chapter 2, “Collecting documentation for specific types of problems,” on page 57 and look for a more likely problem symptom. If you cannot find a more likely symptom, go to “Reporting the problem to IBM” on page 103.

Logon problems

This information provides documentation requirements and diagnosis procedures for logon problems.

The recommended documentation is:

- VTAM full buffer contents trace.

To see the data in the buffer contents trace, set CONFTEXT=NO in the TSOKEY00 member of SYS1.PARMLIB before starting TSO/VTAM.
- VTAM internal trace with MODE=EXT and OPTION=(API, MSG, NRM, PIU, SSCP, PSS).

The VTAM internal trace may not be required. Review the diagnosis procedure for your problem to see whether it is required.

Note: API, MSG, NRM, PIU, and SSCP are always running internally.

1. Did your first logon using USS commands fail?

If so, continue with the next step.

Otherwise, go to step 5 on page 107.

2. Under the information about unformatted system services (USS) tables in the z/OS Communications Server: SNA Resource Definition Reference, review the process for setting up the USS table and using USS commands. Check for the following errors:

- Is your logon command syntax incorrect?

If so, try to log on using the correct command syntax.

Otherwise, continue with the next step.

- Is the logmode name incorrect?

If the logmode name is specified incorrectly, or if a default logmode entry that is inappropriate for the device type is used, you will get USSMSG7. Look up the accompanying sense code in z/OS Communications Server: IP and SNA Codes and correct the logmode name.

3. Can you log on to TSO without using USS commands?

If you cannot log on at all, go to step 4 on page 107.

If you can log on, start the VTAM buffer contents trace and log on again. Look at the trace output to see what session parameters are contained in the BIND, and compare those parameters to the ones in your logmode table.

If the session parameters in your logmode table are incorrect, make the necessary corrections. Also, make sure that the DLOGMOD operand specifies the correct logmode table entry. (For more information about defining TSO/VTAM session parameters, see *z/OS Communications Server: SNA Network Implementation Guide*.)

If you still cannot identify the problem, go to “Reporting the problem to IBM” on page 103.

4. Are you unable to log on at all?

- If this is your first logon attempt from the device as well as your first logon attempt using USS commands, go to step 5.
- If this is not your first logon attempt from the device, go to “Reporting the problem to IBM” on page 103.

5. Did your first logon from a particular device fail?

If so, continue with the next step.

Otherwise, go to step 7 on page 108.

6. Check for an error in the logmode table, or the MODEENT macro.

(These are described in *z/OS Communications Server: SNA Resource Definition Reference*.)

- a. If you receive message IKT029I with return code X'210000' or X'220000', the BIND has been rejected. The following steps can help you find the portion of the BIND that is not valid:
 - 1) Locate the BINFM in the BIND. BINFM must be X'02' or X'03'. (For more information on coding the BIND, see *z/OS Communications Server: SNA Programming*.)
 - 2) If a PSERVIC is coded, see *z/OS Communications Server: SNA Resource Definition Reference* to make sure that all fields are coded correctly.
- b. Check to see whether the DLOGMOD name on the terminal definition statement is a valid logmode table entry.

If it does not match an entry in the logmode table, the first entry in the logmode table is used as the default. The parameters on the default logmode table entry may not be appropriate for your device type, and as a result, the wrong BIND image may be passed to the logon exit and a CLSDST PASS failure may occur.

Note: You can see this failure in the VTAM internal trace using the API option. For more information on the VTAM internal trace, see *z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT*.

- c. Check the logmode table entry to see whether the parameters are specified correctly for your device type.

If not, the wrong BIND image may be passed to the logon exit and a CLSDST PASS failure may occur.
- d. Check to see whether the MODEENT macro is defined correctly.

If it is not, the terminal may reject the BIND, or the terminal may indicate to the logon exit that the terminal is not supported by TSO/VTAM.
- e. If you have not identified the problem, and if users can log on to TSO from other terminals, start the VTAM buffer contents trace and check the BIND.

- f. If logon attempts fail for all terminals, or if the BIND in the buffer contents trace is what you expected, run the VTAM internal trace with options API, PIU, SSCP, and MSG.
 - g. If you still have not resolved the problem, go to “Reporting the problem to IBM” on page 103.
7. **Did previous logons succeed, but now you cannot log on?**
If so, continue with the next step.
 8. **Did you get message IKJ608I?**
If not, go to step 12.
If so, continue with the next step.
 9. **Is message IKJ608I followed by messages IST804I, IST400I, and IST805I?**
If so, the CLSDST PASS or OPEN ACB may have failed. Continue with the next step.
 10. **Is this a cross-domain logon?**
If not, go to step 11.
If this is a cross-domain logon, you may have a VTAM definition problem. For more information on defining TSO/VTAM and logical units that must access TSO/VTAM in a cross-domain environment, see *z/OS Communications Server: SNA Network Implementation Guide*.
 11. **Is there a TSO/VTAM APPLID that is not valid?** (For more information on defining APPLIDs, see *z/OS Communications Server: SNA Network Implementation Guide*.)
If so, correct the APPLID. After you have corrected the APPLID, you must deactivate the application and reactivate it to cause VTAM to reload the correct APPLID. This should fix the problem.
If not, continue with step 12.
 12. **Did you get message IKT111I?**
If not, go to step 13.
If so, check the message text to determine the reason for the logon failure. See *z/OS Communications Server: SNA Messages* for additional diagnostic information for particular messages.
 13. **Is an ABEND0AB with return code X'0105' or X'0203' associated with the logon attempt?**
If not, go to step 14 on page 109.
If so, check the LOGREC entry for the additional information shown in Table 10.

Table 10. ABEND0AB information in LOGREC

Reason code	Register	Contents
X'0105'	6	RPLRTNCD, RPLFDB2, and RPLDAF
	7	RPLFDBK2 (the word of sense)
X'0201'	8	ACBERFLG (for OPEN ACB failure)
	9, 10	TVWA ACB Name (TSOnnnn)
X'0202'	8	ACBERFLG (for OPEN ACB failure)
	9, 10	TVWA ACB Name (TSOnnnn)
X'0203'	5	ACBERFLG (for CLOSE ACB failure)

Table 10. ABEND0AB information in LOGREC (continued)

Reason code	Register	Contents
	6, 7	TVWA ACB Name (TSOnnnn)

Then go to “Reporting the problem to IBM” on page 103.

- If none of the previous situations apply, start the VTAM buffer contents trace and the VTAM internal trace, and trace the logon attempt. Then go to “Reporting the problem to IBM” on page 103.

TSO/VTAM abends

TSO/VTAM issues several unique abends. This information briefly describes the causes and documentation requirements for each one. Use the information provided here and in *z/OS Communications Server: IP and SNA Codes* to try to resolve the problem. If you are not able to do so, go to “Reporting the problem to IBM” on page 103.

ABEND0AB

ABEND0AB occurs when a VTIOC module issues a VTAM macroinstruction that fails. Depending on the values of the RPLRTNCD and RPLFDB2 fields, the macro may be retried. If the retry fails, ABEND0AB is issued.

Table 11. ABEND0AB information in a dump of SDWA

Offset	Length (bytes)	Description
X'280'	8	Terminal name
X'289'	1	RPL request type: X'22' = SEND; X'23' = RECEIVE
X'28A'	1	RPLRTNCD
X'28B'	1	RPLFDB2
X'28C'	4	RPLFDBK2 (Sense code)

If this happens during the execution of a SEND or RECEIVE, the session is placed in reconnect status.

The recommended documentation is:

- Contents of register 15.
This contains the reason code, which is explained in *z/OS Communications Server: IP and SNA Codes*.
- Message text for message IKT116I.
- The software LOGREC entry.
If you have a LOGREC entry, look at an unformatted dump of the SDWA. Table 11 describes the pertinent data you should look for in the dump.
- The dump that is created automatically for this abend.
For abends associated with I/O errors, a dump is not generated automatically unless the RCFBDUMP parameter of the TSOKEY00 member of SYS1.PARMLIB is set for it.

- For errors that occur during session initialization or termination, run the VTAM internal trace with the options API, PIU, MSG, SSCP, and PSS.

Note: The options API, MSG, PIU, and SSCP are always running internally, but you may want to run the VIT with MODE=EXT to be certain that you get the expected output.

ABEND0AC

ABEND0AC occurs when an error halts TCAS processing.

The recommended documentation is:

- Contents of register 15.
This contains the reason code, which is explained in z/OS Communications Server: IP and SNA Codes.
- TWARSON (IKTCASWA + X'02').
This also contains a reason code. See z/OS Communications Server: IP and SNA Codes or see Table 48 on page 649 to determine what document contains the MVS system codes.

ABEND0AD

ABEND0AD occurs when the TSO/VTAM queue manager has a problem manipulating storage for the input and output queues.

The recommended documentation is:

- Contents of register 15.
This contains the reason code, which is explained in z/OS Communications Server: IP and SNA Codes.
- The dump that is created automatically for this abend.

ABEND15D

ABEND15D occurs when the issuer of a TGET, TPUT, or TPG macro passes a data area that is not valid to the SVC 93 modules. A TPUT or TPG request requires read access to the area, and a TGET requires write access. An ABEND0C4 occurs when IKTVTPUT or IKTVTGET tries to validate the data areas passed from the application program, and IKT93EST changes the ABEND0C4 to an ABEND15D.

The recommended documentation is:

- SLIP dump of the ABEND0C4. To determine what document describes the SLIP dump, see Table 48 on page 649.
- GTF trace of SVC 93 entries.

Note: Either of these will show the address that is not valid.

Parameter initialization problems

This topic provides information on finding the parameter member containing the TSO/VTAM initialization parameters.

To find the member name for the initialization parameters:

- Check the system log for the TSO start command issued. The MEMBER or MBR option may have been used to specify the member name. The MEMBER option overrides any other methods of specifying the member name.
- Check the TSO start procedure for any variables used, such as MBR, for the member name.
- The default member name is TSOKEY00.

The parameter member can be found in:

- The data set defined by the PARMLIB DD statement in the TSO start procedure.
- A data set in the logical parmlib concatenation (for z/OS). Refer to z/OS MVS System Commands for information on using the MVS DISPLAY PARMLIB command to display information about the logical parmlib setup for a system.
- SYS1.PARMLIB.

Hung terminal problems

This information provides documentation requirements and diagnosis procedures for problems with hung terminals. Use this information if this problem occurs while you are using TSO/VTAM. This procedure helps you determine when the hang occurred and what was happening at that time.

1. If the problem occurs during logon or logoff, get the recommended documentation and go to “Reporting the problem to IBM” on page 103.

The recommended documentation is:

- VTAM full buffer contents trace.

To see the data in the buffer contents trace, set CONFTXT=NO in the TSOKEY00 member of SYS1.PARMLIB before starting TSO/VTAM.

- GTF trace of SVC 93 and SVC 94 entries. See Table 48 on page 649 to determine what document describes the SVC 93 and SVC 94 entries.
- TGET/TPUT trace.

The TGET/TPUT trace creates trace entries for all TGET/TPUT/TPG data except address space ID TPUTs.

You can get the TGET/TPUT trace by issuing the MODIFY TRACE command with TYPE=TSO. The MODIFY command is described in z/OS Communications Server: SNA Operation.

- Dump of the nucleus, CSA, and user's address space.

If you think you might not be able to re-create the problem, take a dump *before* you try to clear the hang.

If input (such as ATTN, RESET, or ENTER) clears the hang, both a dump and traces may be necessary. You may want to start the traces, take a dump with the terminal hung, clear the hang, and then stop the traces.

- Collect additional general information:
 - Try to determine whether the error is related to a certain type of hardware or a certain protocol (SNA or non-SNA).
 - Try to determine whether the hang is related to a particular application program or type of application program (full-screen, graphics, and so on).

If so, do other similar types of applications also hang?

2. Was the last data that was sent from the application program to VTAM sent to the terminal before the hang occurred?

To determine this, look for a VTAM buffer contents trace entry that corresponds to the last TPUT trace entry.

- If you see these corresponding entries, the output was sent to the terminal.
3. If the keyboard locked after data was sent to the terminal, check the outbound buffer contents trace entry for:
 - Bracketing indicators in RH byte 2
 - Change direction indicator in RH byte 2
 - Write control character to unlock the keyboard in byte 2 of the output request unit

If an end bracket was sent, the keyboard should be available.
 If a change direction was sent and the keyboard has been unlocked, the keyboard should be available.
 If a TGET is issued after a full-screen TPUT, TSO/VTAM should unlock the keyboard.
 4. If the keyboard has not been unlocked, see what TPUT was issued last. (The option flag bytes in the TPUT entry show what TPUT it is.)
 - For a NOEDIT or TPG TPUT, TSO/VTAM will not unlock the keyboard. The application program is supposed to send a write control character to unlock the keyboard.
 - For other TPUT options, determine whether a TGET is outstanding.
 If a TGET is outstanding, TSO/VTAM should unlock the keyboard.
 If no TGETs are outstanding, contact the group responsible for the application program.
 5. If the last activity before the hang was input from the terminal, was the data passed to the application program?(If it was, the TGET trace entry corresponds to the inbound VTAM buffer contents trace entry.)
 - If so, determine whether the application program ever issued another TPUT.
 If the application program never issued another TPUT, contact the group responsible for the application program.
 - If data has been received by TSO/VTAM, but it has not been sent to the terminal, go to “Reporting the problem to IBM” on page 103.
 6. If you have not identified the problem, go to “Reporting the problem to IBM” on page 103.

Incorrect output problems

Two main types of incorrect output problems are discussed in this information: “Screen management problems” and “Screen size problems” on page 117. Screen management problems involve mode errors, exception responses, and problems with the data on the screen. Screen size problems involve an incorrect or unexpected screen size, either in a particular mode or all the time. Choose the one that is most like your symptoms, and follow the procedure for that problem.

Screen management problems

This information provides the documentation requirements and diagnostic procedures for problems displaying data on the screen. This information deals with five types of screen management problems:

- Function error (incorrect screen management for mode).
- Exception responses.
- Extra or missing data.
- Data is not placed correctly on the screen.

- Data appears to be translated incorrectly.

Choose the one that most closely matches your symptoms and follow the procedure for that problem.

Note: Problems with incorrect screen sizes are addressed in “Screen size problems” on page 117.

The recommended documentation is:

- VTAM full buffer contents trace.

To see the data in the buffer contents trace, set CONFTXT=NO in the TSOKEY00 member of SYS1.PARMLIB before starting TSO/VTAM.

- GTF trace of SVC 93 and SVC 94 entries. See Table 48 on page 649 to determine what document describes the SVC 93 and SVC 94 entries.
- TGET/TPUT trace.

The TGET/TPUT trace creates trace entries for all TGET/TPUT/TPG data except address space ID TPUTs.

You can get the TGET/TPUT trace by issuing the MODIFY TRACE command with TYPE=TSO. The MODIFY command is described in z/OS Communications Server: SNA Operation.

- Dump of CSA storage and the user's address space.
This is required only for address space ID TPUT errors.
- May require a full PIU trace or CCW trace with data option.

Function error

If the screen does not function properly for the current mode of operation, do the following steps:

1. Review the SVC 94 entries to determine the mode. The STFSMODE and STLINENO macros set full-screen mode on and off.

- Full-screen mode.

In full-screen mode, the application handles screen management.

If the problem relates to full-screen processing, review the information on full-screen mode. See Table 48 on page 649 to determine what document describes full-screen mode. If this does not describe the full-screen processing that you are experiencing, note the differences.

- Line mode.

In line mode, TSO/VTAM handles screen management.

In line mode, the data generated by the application program is placed line by line down the screen. READY appears on the line below the data, and the cursor appears on the line below that. When the screen is full, TSO/VTAM sends a page prompt to the screen. When you press the ENTER key, TSO/VTAM clears the screen and sends any remaining data to the screen.

If you enter data on the next to the last line of the screen, no page prompt is sent. Instead, TSO/VTAM clears the screen and reshows the data (or command) at the top of the screen.

If this does not describe the line mode processing that you are experiencing, note the differences.

2. Report to the group (TSO/VTAM or the application program) that appears to be responsible for the incorrect screen management.

- If you have access to IBMLink, search for known problems in this area. If no applicable matches are found, report the problem to IBM by using the electronic technical report (ETR) option on IBMLink.
- If you do not have access to IBMLink, call the IBM Support Center at 1-800-IBM-SERV.

Exception responses

Follow these steps for exception responses.

1. Determine if the error indicator reflects an error in the RU portion of the PIU (sense = 1003 or 1005).
If so, continue with this step.
Otherwise, go to step 2.
 - a. See what TPUT was issued. (This is shown in the flag byte of the TPUT trace entry.)
 - b. For a NOEDIT TPUT, TSO/VTAM should not change the data provided by the application program. Compare the data in the TPUT trace, which starts at X'2C' into the trace record, with the VTAM buffer contents trace.
If TSO/VTAM has not changed the data, contact the group responsible for the application program that issued the TPUT.
 - c. For a full-screen TPUT, determine whether the data that is causing the error was generated by the application program. (Data generated by the application program is present in the TPUT trace entry.)
If so, contact the group responsible for the application program that issued the TPUT.
If not, go to "Reporting the problem to IBM" on page 103.
2. Look at the error sense code.
If it is X'800A', the PIU is too long. This is probably a definition error. In this case:
 - a. Look in the TH portion of the buffer contents trace entry to find out the length of the PIU that caused the error.
 - b. See how MAXDATA is defined on the PCCU definition statement in the NCP definition deck. The MAXDATA value that you code should be as large as the largest PIU that is sent to the terminal by an application program.

Note: If you increase the value of MAXDATA, this value should not exceed the product of the MAXBFRU and UNITSZ operands. See z/OS Communications Server: SNA Network Implementation Guide for more information about defining the MAXDATA, MAXBFRU, and UNITSZ operands.
If you have more than one PCCU definition statement, check to see that the right one is being used.
3. If there are other error indicators, get the documentation shown in step 5 on page 93 and continue with that procedure.
4. If you have not resolved the problem, go to "Reporting the problem to IBM" on page 103.

Extra or missing data

Follow these steps if you have an extra or missing data condition.

1. Is the problem with input (data received by the application)? Is too much data being passed to the application?
 - If previous input is added to the end of the current input, the modified data tags may have been set improperly in previous TPUTs. Go to “Reporting the problem to IBM” on page 103.
 - If the backspace key, character delete key, or line delete key is not functioning properly, look for SVC 94 entries that may have changed the default in the STCC macro.

If this does not account for the problem, get a dump of the nucleus, the CSA, and the user's address space. Then go to “Reporting the problem to IBM” on page 103.

Is insufficient data being sent to the application program?

- Find out which TGET option was used. TSO/VTAM edits data sent from the terminal before it passes it to the application program. The type of editing that TSO/VTAM does depends on the TGET option. Certain characters may be deleted, such as control characters, aid characters, and set buffer address (SBA) sequences. Look at the flag bytes in the TGET trace entry to see which TGET option was specified. See Table 48 on page 649 to determine what document describes TGET options.
 - Find out whether any user edit exits are involved. User edit exits are listed in z/OS Communications Server: SNA Customization.
2. Is the problem with output (data sent by the application)? Was extra data sent to the screen?

TSO/VTAM should not generate any printable data. Compare the data portion of the TGET/TPUT trace with the data in the VTAM buffer contents trace.

If you see printable data in the buffer contents trace that is not in the TGET/TPUT trace, go to “Reporting the problem to IBM” on page 103.

If you see the same data in both traces, contact the group responsible for the application program that issued the TPUT. Was data from the application lost?

TSO/VTAM does not generally delete data sent by the application program unless it is doing reshew processing. In reshew processing, TSO/VTAM deletes the full-screen TPUT and sends a reshew character (X'6E') to the full-screen application program. This is shown in the TGET trace entry. Compare the data portion of the TPUT trace entry with the data in the VTAM buffer contents trace entry to see whether any data has been deleted. Determine whether this is a TSO/VTAM problem or an application program error and contact the appropriate group.

See Table 48 on page 649 to determine what document contains more information on reshew processing.

3. If you have not resolved the problem, go to “Reporting the problem to IBM” on page 103.

Data is misplaced on the screen or page

Follow these steps if data is misplaced on your screen or page.

1. Does data wrap around the screen? (Wrapping means that data fills the line and splits inappropriately between lines.)

If it does, continue with this step.

Otherwise, go to step 2 on page 116.

 - a. Find out which TPUT option was issued. Bytes X'12' and X'13' of the TPUT trace entry contain the option flags.

- b. For a NOEDIT TPUT, TSO/VTAM does not edit data, and therefore does not change any SBA sequences that may be issued by the application program. Contact the group responsible for the application program that issued the TPUT.
 - c. For a full-screen TPUT, TSO/VTAM does not generate SBA sequences to place the data on the screen.
If the symptom is incorrect screen size, go to “Screen size problems” on page 117. Otherwise, contact the group responsible for the application program that issued the TPUT.
 - d. If the data was sent without a full-screen or NOEDIT option, go to “Reporting the problem to IBM” on page 103.
2. If the problem is not on a display terminal, go to “Reporting the problem to IBM” on page 103.
 3. Are SBA sequences correct? Check the type of TPUT that was issued. Bytes X'12' and X'13' of the TPUT trace entry contain the option flags. For full-screen and NOEDIT TPUTs, the application program usually generates the SBA sequences that determine where data is placed on the screen.
 4. Is the buffer address incorrect?
 - If the buffer address is not valid for the screen size, contact the group responsible for the application program.
 - If the buffer address is valid for the terminal in its present screen size, go to “Reporting the problem to IBM” on page 103.
 5. If this is the first nonfull-screen TPUT following a full-screen TPUT, look for an SVC 94 trace entry for STLINENO. This macro may be issued by a full-screen application to indicate which line the next nonfull-screen data should appear on.
If the data was placed by the STLINENO macro, contact the support group for the application program that issued the macro.
 6. If you have not resolved the problem, go to “Reporting the problem to IBM” on page 103.

Data appears to be translated incorrectly

Incorrect output is the main symptom of this problem. Sometimes the incorrect output is colons. The problem is probably related to the TPUT option that was specified.

1. Look at the option flag bytes (X'12' and X'13') in the TPUT trace entry to determine what options were used.
2. Determine what editing occurs for each type of TPUT option. See Table 48 on page 649 to determine what document describes editing done by TPUT options.
3. If the incorrect output consists of colons, determine what data from the TPUT trace entry is being edited into the printable character X'7A', a colon.
4. See whether TSO/VTAM is editing correctly.

Note: Many applications use TPUT options that do extensive editing and translation. This allows many different hardware devices to communicate with the application program without causing I/O errors. You may need to write your own user edit exits to make sure that all characters that are valid for your terminals appear on the screen, especially if you are using type 1 logical unit devices.

5. If the editing does not appear to conform to the options specified, check for user edit exits or translation tables. (See *z/OS Communications Server: SNA Customization* for more information on these exits.)
If you have exits or translation tables, verify that they are not causing the problem.
6. Did the application program send incorrect data to VTAM? The data is shown in the data portion of the TPUT trace entry.
If so, the problem is in the application program.
7. If you have not resolved the problem, go to “Reporting the problem to IBM” on page 103.

Screen size problems

This information provides documentation requirements and diagnostic procedures for screen size problems. Two major types of errors occur:

- The screen never operates in the expected size.
- The screen is not always the expected size when you change modes.

Choose the one that most closely matches your symptoms and follow the procedure for that problem.

Recommended documentation includes the following list.

Note: All of the traces may not be required. Read the diagnostic procedure before you get them.

- VTAM buffer contents trace.
To see the data in the buffer contents trace, set CONFTXT=NO in the TSOKEY00 member of SYS1.PARMLIB before starting TSO/VTAM.
- GTF trace of SVC 93 and SVC 94 entries. See Table 48 on page 649 to determine what document describes the SVC 93 and SVC 94 entries.
- TGET/TPUT trace.
The TGET/TPUT trace records all TGET/TPUT/TPG data except address space ID TPUTs.
You can get the TGET/TPUT trace by issuing the MODIFY TRACE command with TYPE=TSO. The MODIFY command is described in *z/OS Communications Server: SNA Operation*.
- May require a full PIU trace or CCW trace with data option.

Screen is never the expected size

This is probably a definition problem.

1. If you are using a USS command to log on, try logging on without it.
If this corrects the problem, review the use of the USS command in *z/OS Communications Server: SNA Resource Definition Reference* and check the following items:
 - Is the terminal a non-SNA 3270?
If it is, does the USS command include a USSPARM macro for the logmode?
If so, VTAM ignores the logmode name from the terminal definition statements and uses its own default BIND image instead.
 - Does the USS command establish a default logmode name?

- If so, the default name overrides the name in the terminal definition statement.
2. If this is not a USS command problem, check to see whether the PSERVIC operand of the MODEENT macro is coded correctly. z/OS Communications Server: SNA Resource Definition Reference explains how to do this.
 - Are primary and alternate sizes coded correctly?
 - For screen switching, is BINPRESZ coded correctly as X'7F'?
 3. If you have not identified the problem, look at the BIND that is sent. This is shown in a VTAM buffer contents trace of the logon.
 - If the BIND image is not what you expected, check the LU definition statement for an incorrect MODETAB or DLOGMOD parameter.
 - If no logmode table or DLOGMOD operand is specified, no PSERVIC is passed to the TSO/VTAM logon exit routine. In this case, TSO/VTAM issues an INQUIRE DEVCHAR macro and VTAM indicates that the terminal is a logical unit. TSO/VTAM then uses the SCRSIZE operand found in TSOKEY00. The default value for SCRSIZE is 480 (12 rows and 40 columns).
 4. If you have not resolved the problem, go to "Reporting the problem to IBM" on page 103.

The screen is not the expected size for the mode

Full-screen mode

The application program controls screen management in full-screen mode. The primary (small) screen size is considered by TSO/VTAM as the default size. The application program can control screen size by sending write commands in TPUTs that it issues. The write commands are X'F5', erase write, and X'7E', erase write alternate. The application program issues X'F5', erase write, to set the primary screen size, or X'7E', erase write alternate, to set the alternate (large) screen size.

If neither command is issued, the screen remains the size it is when the application program enters full-screen mode.

Line mode

TSO/VTAM controls screen management in line mode. It generally uses the large (alternate) screen size when processing TPUTs in line mode. You can use the TSO TERMINAL command STSIZE macro during a session to change the screen size for nonfull-screen processing.

Using output from the VTAM full buffer contents trace, the TGET/TPUT trace for TSO/VTAM, and the GTF trace of the SVC 93 and SVC 94 entries, try to locate the source of the problem.

1. Check the SVC 94 trace entries to see whether the processing is in full-screen mode or line mode. The STFSMODE and STLINENO macros set these modes on and off.

If the processing is in line mode, go to step 6 on page 119. For full-screen mode, continue with step 2.

2. Note if the incorrect screen size is related to entering or exiting full-screen mode.
3. Locate the TPUT trace entry for the data that appears on the screen when the screen is the wrong size. Determine the TPUT options for this TPUT and the one that precedes it by looking at the option flag bytes.

If either is a full-screen TPUT, look at the first data byte.

If the first data byte is an escape character (X'27'), the write command that follows has been specified by the application program. This write command should determine the screen size.

If the write command is different in the buffer contents trace, go to "Reporting the problem to IBM" on page 103.

4. If the first data byte is not an escape character, determine whether a write command (X'F1', X'7E', or X'F5') is provided.

If one of these write commands exists, continue with this step.

If not, go to the next step.

If this write command is different in the VTAM buffer contents trace, go to "Reporting the problem to IBM" on page 103.

If the same write command is displayed in both traces, contact the group responsible for the application program.

5. If a write command is not provided in the TPUT data, and processing is in full-screen mode, determine if the write command generated by TSO/VTAM set the same screen size as the last write command provided by the full-screen application program. To determine this, compare the write command in the buffer contents trace entry with the last one provided in a TPUT trace entry. The write command is located in the data portion of the TPUT trace, at X'2C' into the entry.

If the write commands are the same, go to "Reporting the problem to IBM" on page 103.

If the write commands are different, contact the group responsible for the application program.

6. If processing is not in full-screen mode, determine whether the STSIZE macro set the screen size. To determine this, look in the SVC 94 trace entries. An entry code of X'0A' in the high-order byte of register 0 indicates that the STSIZE macro set the screen size.

If the screen size was set by the application program, contact the group responsible for the application program.

7. If you have not resolved the problem, go to "Reporting the problem to IBM" on page 103.

Performance problems

This information provides documentation requirements and diagnosis procedures for performance problems. Use this information along with the information in "Performance problem" on page 94. In addition to the documentation required in "Performance problem" on page 94, you will need a GTF trace of SVC 93 and SVC 94 entries. See Table 48 on page 649 to determine what document describes the SVC 93 and SVC 94 entries.

1. Are wait and hold options slowing response time?

The application program may be issuing TPUTs and TGETs with the wait or hold options. The wait option indicates that the application program should not regain control until output data has been placed on the output queue (TPUT) or input is available (TGET). A TPUT with a hold option indicates that control should not be returned to the application until the data has reached the terminal. These options may be necessary for screen management, but they prolong response time and increase the number of times the address space is swapped.

2. Is an external system resource slowing response time?

If all users must access the same resource, such as a system catalog, performance deteriorates. This problem is especially severe when exclusive ENQs are used to control access to the resource. To improve performance, redistribute resources.

3. Are high and low buffer extents set at inappropriate values?

If the high and low buffer extents are too close together, output wait states occur. Buffer extents are specified in TSOKEY00, a member of SYS1.PARMLIB. Tune the values of the high and low buffer extents to get optimum performance.

4. Are APPL definition statements coded correctly?

Code the AUTH=NVPACE operand on all APPL definition statements for TSO/VTAM.

If you do not set NVPACE, VTAM indicates that it has already received input data, instead of queuing the response until it receives the input data. Also, the swap count is incremented by two every time the ENTER key is pressed.

5. Are pacing values set correctly for local SNA terminals?

VTAM ignores the NVPACE operand for sessions with logical units in a local major node. Therefore, you must set nonzero pacing values for these logical units.

6. Is the MVS performance group specified correctly?

Set the application program's performance group approximately five to ten percent lower than the VTAM performance group. To see the application program's performance group specification, look at the dispatching priority in the task's TCB.

7. If you have not resolved the performance problem, go to "Performance problem" on page 94.

Part 2. Diagnostic procedures

Chapter 4. Using DISPLAY and MODIFY operator commands

You can control and monitor the VTAM program network with the following start options and operator commands:

- DISPLAY commands
 - “Using VTAM DISPLAY commands for problem determination” on page 124
 - “Display buffer pool use” on page 127
 - “DISPLAY CSDUMP” on page 128
 - “DISPLAY CSMUSE” on page 128
 - “DISPLAY EE” on page 128
 - “DISPLAY EEDIAG” on page 131
 - “Display Enterprise Extender connection network unreachable partner information” on page 149
 - “Display HPR route test” on page 150
 - “Display ID for an RTP connection” on page 152
 - “Display ID for an RTP PU with HPRDIAG=YES” on page 152
 - “Display ID for an RTP PU with HPRDIAG=YES and CLEAR=ALL” on page 154
 - “Display NCP storage” on page 155
 - “Display path tables” on page 155
 - “Display resource status” on page 156
 - “Display resources in a pending state” on page 156
 - “Display route status” on page 156
 - “Display route test” on page 163
 - “Display RTPS options” on page 166
 - “Display TDU information” on page 167
 - “Display traces” on page 175
 - “Display VTAM storage” on page 175
 - “Display workload information for a device” on page 175
- MODIFY commands
 - “Using VTAM MODIFY commands for problem determination” on page 177
 - “Issuing the MODIFY CSDUMP command” on page 177
 - “Modifying input/output problem determination” on page 177
 - “Modifying message module identification” on page 178
 - “Modifying NCP intensive mode recording” on page 179
 - “Modifying SDLC link level 2 test” on page 179
 - “Issuing the MODIFY TOPO command to clear EE connection network unreachable partner information” on page 180
 - “Modifying tuning statistics” on page 181
 - “Issuing the MODIFY VTAMOPTS command to change start option values” on page 181

For information about VTAM start options, see *z/OS Communications Server: SNA Resource Definition Reference*.

Note: You can also use the NetView program to monitor and collect error statistics from the VTAM network.

Using VTAM DISPLAY commands for problem determination

VTAM provides DISPLAY (D) commands to show status and other information about network resources. The following list shows what information is displayed for each of the VTAM DISPLAY commands. For more information about the syntax and output of these commands, see *z/OS Communications Server: SNA Operation*.

Command

Information Displayed

D ADJCLUST

Adjacent cluster table definitions in the current ADJCLUST table

D ADJCP

Status of adjacent CP major nodes

D ADJSSCPS

Adjacent SSCP tables

D APING

Existence of route to LU 6.2 resource; route information; throughput statistics for conversation on route

D APINGDTP

Number of APINGD transaction programs to run concurrently

D APINGTP

The number of APING command transaction programs permitted to run concurrently for sending APING requests to other nodes

D APPLS

Status of application program major and minor nodes

D APPNTOSA

APPN to subarea CoS mapping table

D AUTOLOG

Information about controlling applications that have pending autologon requests

D BFRUSE

VTAM buffer usage

D BNCOSMAP

Native and nonnative CoS mapping defined for a border node

D CDRMS

Status of cross-domain resource manager major and minor nodes

D CDRSCS

Status of cross-domain resources (including independent LUs)

D CLSTRS

Status of clusters (PUs in NCP, local SNA, and switched major nodes)

D CNOS

Change-number-of-sessions characteristics for LU 6.2 application programs

D CONVID

Conversations with LU 6.2 application programs

- D COS**
Class of Service table information
- D CPCP**
CP-CP session status
- D CSDUMP**
Current[®] dump triggers that are set by the MODIFY CSDUMP command or by the CSDUMP start option
- D CSM**
Information on the use of storage managed by the communications storage manager (CSM)
- D CSMUSE**
Displays the detail usage of storage managed by the communications storage manager (CSM) for one or more storage pools.
- D DIRECTORY**
Information maintained by central directory server
- D DISK**
Disk contents of 3720 or 3745 Communication Controller
- D DLURS**
All dependent LU requesters for which this host acts as Dependent LU Server
- D EE** Information about Enterprise Extender connections
- D EEDIAG**
Display various diagnostic information about one or more EE connections.
- D EXIT**
Status of user-written exit routines
- D GRAFFIN**
Affinity information for generic resources
- D GROUPS**
Status of line groups
- D GRPREFS**
The contents of the generic resource preference table
- D ID** Individual major or minor nodes
- D INOPCODE**
The attributes for every INOPCODE defined to VTAM or for every INOPCODE defined within a single VTAM module
- D INOPDUMP**
The global status for INOPDUMP
- D LINES**
Status of lines and channel links
- D LMTBL**
LU-mode table for LU 6.2 application programs
- D LUGROUPS**
LUGROUP major nodes, model LU groups, and model LUs
- D MAJNODES**
Status of major nodes

- D MODELS**
Model PUs and LUs
- D NCPSTOR**
Storage contents of 3720 or 3745 Communication Controller
- D NETSRVR**
Network node server information
- D PATHS**
Dial-out path information
- D PATHTAB**
Status of explicit routes and virtual routes
- D PENDING**
Resources in a pending state
- D ROUTE**
Status of explicit routes and virtual routes; existence of routes; whether a route is operational; whether a route is blocked
- D RSCLIST**
Resources whose names match a particular pattern
- D RTPS**
Information concerning HPR RTP connections
- D SAMAP**
Status of subarea mapping
- D SATOAPPN**
Subarea-to-APPN Class of Service mapping table
- D SESSIONS**
Session status information
- D SNSFILTR**
Current active SAW sense filter
- D SRCHINFO**
Information about outstanding subarea and APPN search requests
- D STATIONS**
Status of cross-subarea link stations
- D STATS**
Storage information for use with the storage estimate worksheets appendix of z/OS Communications Server: New Function Summary.
- D STORUSE**
Storage usage for storage pools and data spaces
- D TABLE**
Table type, use count, and users
- D TERMS**
Status of device-type LUs (terminals)
- D TGPS**
Transmission group profiles
- D TNSTAT**
Tuning statistics information

D TOPO

Topology of APPN network (information about nodes and transmission groups)

D TRACES

Status of VTAM and NCP traces

D TRL

Information about the TRL major node or about a single TRLE definition statement

D TSOUSER

Status of a TSO user ID

D USERVAR

USERVARs and the application programs associated with them

D VTAMOPTS

Start options

D VTAMSTOR

Display storage contents associated with a storage address

Display buffer pool use

You can use the DISPLAY BFRUSE command to display information about buffer use. In response to this command, VTAM indicates that the display is for buffer use and issues a series of messages that contain monitoring information. For each buffer pool, this information includes:

- Buffer pool ID
- Flags (Q or F): Q shows that a request is queued for this pool; F shows that dynamic buffering has failed for this pool
- Size of each buffer in this pool
- Current total number of buffers in this pool
- Current count of buffers available (the number not in use)
- Largest number of buffers this pool has expanded to at any time
- Largest number of buffers in use at any time
- Cumulative count of the number of times each buffer pool has expanded
- Expansion and contraction thresholds
- The expansion increment (the number of buffers to be added to a buffer pool during dynamic expansion)
- VTAM intermediate routing node buffer use limit (IRNLIMIT), current buffer use, and maximum buffer use
- VTAM CSA buffer use limit (CSALIMIT), current buffer use, and maximum buffer use
- Maximum amount of CSA in use since VTAM was started
- Current amount of VTAM private storage and maximum amount of VTAM private storage

If the DISPLAY BFRUSE command is used while an SMS (buffer use) trace is running, the fields MAX TOTAL, MAX USED, and TIMES EXP reflect buffer usage only since the last trace record was written, because the SMS trace resets these fields. For more information about the syntax and output of the DISPLAY BFRUSE command, see z/OS Communications Server: SNA Operation.

DISPLAY BFRUSE output can help you identify possible sources of problems. The following chart shows some problem symptoms and the corresponding buffers to check in SMS trace output:

For this symptom:	Check this buffer pool:
I/O hang	IOBUF
Session failure	CRPLBUF and LPBUF
VTAM hang	LPBUF

Storage problems can also be related to an I/O device. For further information, see “Display workload information for a device” on page 175.

DISPLAY CSDUMP

You can use the DISPLAY CSDUMP command to display the current dump triggers set by the MODIFY CSDUMP command or by the CSDUMP start option. The display shows the current CSDUMP message and sense code triggers that will initiate a dump. If either the message or the sense code trigger does not exist, then NONE is indicated.

See the DISPLAY CSDUMP and MODIFY CSDUMP commands in z/OS Communications Server: SNA Operation for more information.

DISPLAY CSMUSE

You can use the DISPLAY CSMUSE command to determine the CSM (communications storage manager) managed storage growth used by the components of z/OS Communications Server. The DISPLAY CSMUSE command allows IBM service to evaluate the use of storage managed by the CSM. Although the command is similar to DISPLAY CSM command, it provides a lower level of detail regarding storage usage. Therefore, the output of this command is different from that of DISPLAY CSM.

See the DISPLAY CSMUSE command in z/OS Communications Server: SNA Operation for more information.

See the description of monitor IDs in z/OS Communications Server: IP and SNA Codes for more information.

DISPLAY EE

You can use the DISPLAY EE command to obtain information about Enterprise Extender. This command has various formats providing general Enterprise Extender information as well as detailed connection throughput statistics. A few of the display command formats will be shown along with some of the important messages.

- To display general Enterprise Extender information in summary format, use the following command:

```
D NET,EE
```

- Message IST1685I identifies the job name of the TCP/IP stack which Enterprise Extender is using.

- Message IST2004I displays the Enterprise Extender Logical Data Link Control (LDLC) timer and disconnect timer values associated with the PORT definition statement.
- Message IST2005I shows the number of seconds VTAM waits for name-to-address resolution requests to complete before canceling the request. The value displayed is associated with the Enterprise Extender port and affects only local HOSTNAME name-to-address resolution requests. When an EE line is in the process of being activated, and VTAM is performing name-to-address resolution for the local HOSTNAME, a display of the line (D NET,ID=*linename*) will show a state of PGAIN (Pending GetAddrInfo). If an EE line is hung in PGAIN state, you can perform the following steps to identify why the local HOSTNAME name-to-address resolution is not completing:
 1. Verify that the TCP/IP stack identified in message IST1685I is active.
 2. Verify that the TCP/IP resolver is active.

Result: If IPRESOLV displays the value of 0, VTAM will wait infinitely for the name-to-address resolution to complete. In this situation, VTAM relies on the TCP/IP resolver to time out the resolution. For more information, see *z/OS Communications Server: IP Configuration Guide*.
- IST2008I displays the IP Type of Service (ToS) values associated with each of the Enterprise Extender port priorities.
- IST2021I displays the total number of active Enterprise Extender connections.
- To display general Enterprise Extender information in detail, use the following command:


```
D NET,EE,DET
```

 - The detailed format of the general Enterprise Extender display provides detailed information for each local IP address. Message IST1680I is the first message of a message group. The information for each local IP address is displayed between the IST924I messages.

For each local IP address active to Enterprise Extender, the following information is displayed:

 - Message IST2004I displays the Enterprise Extender Logical Data Link Control (LDLC) timer and disconnect timer values used by this local static VIPA.
 - Message IST1910I/IST1911I displays the local HostName (if applicable).
 - Message IST2009I displays the total number of RTP pipes traversing EE connections associated with this local IP address. This message also displays the total number of LU-LU sessions associated with these RTP pipes.
 - IST2010I displays the number of Enterprise Extender lines which have been INOPed because of SRQRETRY exhaustion. This count is maintained from the time the first EE line (associated with this specific local IP address) is activated, until the last line (associated with this specific local IP address) is deactivated. When the last line is deactivated for this local IP address, this counter will be cleared.
 - For each VRN, the following information will be displayed:
 - Message IST1324I displays the VNNAME and VNGROUP, along with a LOCAL or GLOBAL VRN indicator.
 - Message IST2011I displays the number of available lines associated with this VRN.

Guideline: Verify that IST2011I displays enough lines to support the number of VRN connections that can exist. If the number of available lines drops to zero, new EE connections associated with this connection network will fail to connect.

- IST2012I displays the number of active EE connections associated with this VRN.
- Message IST2013I displays the number of available lines for predefined EE connections associated with this local IP address.

Guideline: Verify that IST2013I displays enough lines to support the number of predefined connections that can exist for this local IP address. If the number of lines is not large enough, new EE connections will fail to connect. If the number of available lines displays as zero, this might mean that all Enterprise Extender lines are associated with the Connection Network (CN) groups. For this case, all available lines that are associated with CN groups are available for predefined connections as well. Lines will be selected from the local CN groups first. If no local CN lines are available, then lines will be selected from the global CN groups.

- To display Enterprise Extender connection information in detail, use the following command:

```
D NET,EE,ID=puname or linename,DET
```

Tip: The DISPLAY EE commands have various formats in which the connection information can be displayed. The example here uses the ID=operand. The DISPLAY EE command with the HOSTNAME/IPADDR operands provides essentially the same information.

- Message IST2022I displays the date and time of the Enterprise Extender connection activation.
- Message IST2114I displays the initial, maximum, and current LIVTIME values for an EE connection.
- Message IST2025I displays the number of LDLC signals over this EE connection which did not receive a response on the first try. The signal required at least one retransmission before a response was received from the EE partner.
- Message IST2026I is closely associated with message IST2025I. The value displayed here indicates the number of LDLC signals over this EE connection which did not receive a response up to SRQRETRY times. It required the signal to be retransmitted SRQRETRY times, at which time a response was received from the EE partner.

Tip: If this display is issued repeatedly over a period of time, and the values displayed in messages IST2025I or IST2026I continue to grow, this indicates that there is most likely a problem in the network. Network congestion is a possible problem which might lead to Enterprise Extender failure. Increasing the LIVTIME, SRQRETRY, and SRQTIME values on the EE XCA PORT macro will allow Enterprise Extender connections to tolerate longer network delays. However, if severe network delays are encountered, it is most likely Enterprise Extender connections will INOP due to timeout conditions.

- Message IST2029I displays the largest MTU size that Enterprise Extender will send over the IP network for this connection. When policy-based routing is in effect, the MTU size might be different for each of the ports, depending on the routes chosen for EE traffic. This message is issued for each of the five EE ports regardless of whether policy-based routing is in effect and regardless of whether the display is for an IPv4 or IPv6 connection. The MTU size (both IPv4 and IPv6) might change during the life of the EE connection. The displayed value is obtained in the following manner: Message IST2029I

displays the largest MTU size that Enterprise Extender will send over the IP network for this connection. When policy-based routing is in effect, the MTU size might be different for each of the ports, depending on the routes chosen for EE traffic. This message is issued for each of the five EE ports regardless of whether policy-based routing is in effect and regardless of whether the display is for an IPv4 or IPv6 connection. The MTU size (both IPv4 and IPv6) might change during the life of the EE connection. The displayed value is obtained in the following manner:

- Initially, VTAM queries the TCP/IP stack for its MTU size and sets the EE connection to use this value. This MTU size has already been reduced to account for various header lengths such as the IP, UDP, and LLC headers necessary for EE traffic.
- VTAM also takes into account the VTAM MTU operand value, if specified. The MTU operand may be specified on three types of VTAM major nodes:
 - For EE connection networks, this parameter may be defined on the connection network GROUP definition statements in the EE XCA major node.
 - For dial-in Enterprise Extender connections which have their associated PUs dynamically created, this parameter may be defined on the model major node (DYNTYPE=EE) PU definition statement.
 - For predefined Enterprise Extender connections, this parameter may be defined on the PU definition statement in the switched major node.
- VTAM then takes the lesser of the TCP/IP stack's computed MTU size and the VTAM defined MTU operand value (if specified). If the TCP/IP stack presents a value less than 768 bytes, VTAM sets the MTU to 768 because this is the smallest packet size allowed by the HPR architecture.
- Generally the MTU size for an EE connection is fairly constant when the EE connection is established. However, in the event the TCP/IP stack's MTU size changes, RTP pipes with endpoints on the same node as the TCP/IP stack dynamically detect these changes when their outbound packets are being transmitted. The MTU size changes because of the following reasons:
 - New IP routes come available with different local MTU sizes
 - Existing IP routes become unavailable.
 - Path MTU discovery is enabled for IPv4 or IPv6 EE connections (See the PMTUD start option for details, z/OS Communications Server: SNA Resource Definition Reference), and path MTU changes are discovered in the IP network.
- Message IST2038I and IST2039I display the number of packets and the number of bytes that have been retransmitted. These counts are displayed for each port priority level. If the values displayed in these two messages increase over time, this indicates problems within the transport network. Large numbers of retransmissions due to network congestions will result in poor RTP performance. If excessive retransmissions occur, RTP path switching might occur.

DISPLAY EEDIAG

The DISPLAY EEDIAG command is used to display diagnostic information about one or more Enterprise Extender connections.

The REXMIT format lists Enterprise Extender connections whose retransmission rate, calculated at each port priority, meet, or exceed a specified threshold. The SRQRETRY format lists Enterprise Extender connections that are experiencing

LDLC signal retries that meet or exceed a specified threshold. A CLEAR function enables the diagnostic counters used by these commands to be cleared for the next measurement interval. A few of the display command formats are shown in this information along with some of the important messages.

Using the REXMIT option on D EEDIAG

Find all Enterprise Extender connections whose retransmission rates meet or exceed 5%, display the output in summary format, and clear all diagnostic counters after command processing is complete:

```
D NET,EEDIAG,REXMIT=5,CLEAR,LIST=SUMMARY
```

See z/OS Communications Server: SNA Operation for the display output.

Tip: The DISPLAY EEDIAG command has various formats in which the connection information can be displayed. This sample command does not use any command filters. The DISPLAY EEDIAG command can be specified with the ID, HOSTNAME, or IPADDR filters to limit the scope of the search.

- Message IST2067I displays the date and time the DISPLAY EEDIAG command was issued.
- Message IST2069I displays the date and time when the REXMIT counters were last cleared. The date and time provided in the message combined with the date and time taken from message IST2067I, provides the time interval in which the retransmission metrics were collected.
- Message IST2036I displays the total number of network layer packets (NLP) that have been sent across this EE connection for this specific priority. This value is maintained from the time and date specified in message IST2069I. Message IST2036I is associated with a specific port priority represented by a subgroup of messages; the subgroup begins with either message IST2030I, IST2031I, IST2032I, IST2033I, IST2034I, or IST2035I.
- Message IST2038I displays the retransmission rate for this EE connection. The retransmission rate is valid from the time and date specified in message IST2069I. If the retransmission rate is excessive, this indicates problems within the IP transport network. Large numbers of retransmissions because of network congestions result in poor RTP performance. If excessive retransmissions occur, RTP path switching might occur. Message IST2038I is associated with a specific port priority represented by a subgroup of messages that begins with either message IST2030I, IST2031I, IST2032I, IST2033I, IST2034I, or IST2035I.

Rule: If you specify REXMIT=xx with the LIST=SUMMARY option, the display provides an overall retransmission rate for all port priorities. It is possible the retransmission rate displayed in message IST2068I, associated with all port priorities, is smaller than the specified REXMIT=xx rate. This means that at least one of the EE port priorities for this EE connection is experiencing a retransmission rate that meets or exceeds the specified rate. In this case, the LIST=SUMMARY option displays the message groups for the specific port priorities that meet or exceed the specified rate, along with the summary of all port priorities.

- Messages IST2071I, IST2072I, and IST2073I all display the number of EE connections that had either the REXMIT counter, SRQRETRY counter, or both counters cleared as part of command processing. The number of EE connections cleared might be larger than the number of EE connections displayed in message IST2042I. For example, a local IPADDR that is used by 500 EE connections might be specified on the command. The CLEAR=REXMIT option will clear the REXMIT counters for all 500 EE connections. However, only one of these EE

connections might be experiencing retransmission problems. In this case, message IST2042I lists only one connection displayed.

Using the SRQRETRY option on D EEDIAG

To locate all Enterprise Extender connections that are experiencing LDLC retries of three or more attempts before receiving a response from the partner EE node, use the command:

```
D NET,EEDIAG,SRQRETRY=3
```

Tip: The DISPLAY EEDIAG command has various formats in which the connection information can be displayed. This sample command does not use any command filters. The DISPLAY EEDIAG command can be specified with the ID, HOSTNAME, or IPADDR filters to limit the scope of the search.

- Message IST2004I displays the Enterprise Extender Logical Data Link Control (LDLC) timer and disconnect timer values used by the local static VIPA.
- Message IST2074I displays the number of times an LDLC TEST command (SRQRETRY attempt) had to be retried before receiving a response from the EE partner. It also displays the number of instances in which this number of retry attempts was required to receive a response from the EE partner.

Tip: The information displayed in the various IST2074I messages can be useful in tuning EE timer operands such as LIVTIME, SRQTIME, and SRQRETRY.

See z/OS Communications Server: SNA Operation for a sample display of DISPLAY EEDIAG using the SRQRETRY option.

How Enterprise Extender times out an inactive connection

During periods of inactivity, no inbound HPR traffic is detected by this EE endpoint for a period of time equal to the LIVTIME value, and an LDLC test request is sent to the EE partner. If no response is received from the EE partner within the SRQTIME interval, another LDLC TEST request is sent. This process is repeated for the number of times specified by SRQRETRY. If no TEST response is received from the EE partner after the last SRQRETRY attempt, the EE connection is disconnected. The format of the IST1430I message that is issued for this scenario is as follows:

```
IST1430I REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT
```

Enterprise Extender inactivity example

Assume the following values:

```
LIVTIME=15  
SRQTIME=15  
SRQRETRY=3
```

If an EE connection is not receiving data from the partner for 15 seconds (LIVTIME), VTAM sends a test frame to test the connection. If the test does not receive a response within 15 seconds (SRQTIME), VTAM repeats this up to three more times (SRQRETRY). If after the third retry attempt no response has been received, the EE connection is disconnected. VTAM uses the following formula: $LIVTIME + (SRQTIME * (SRQRETRY + 1))$. In this example, it would take roughly 75 seconds to disconnect the EE connection. Figure 17 on page 134 shows the Enterprise Extender inactivity flows.

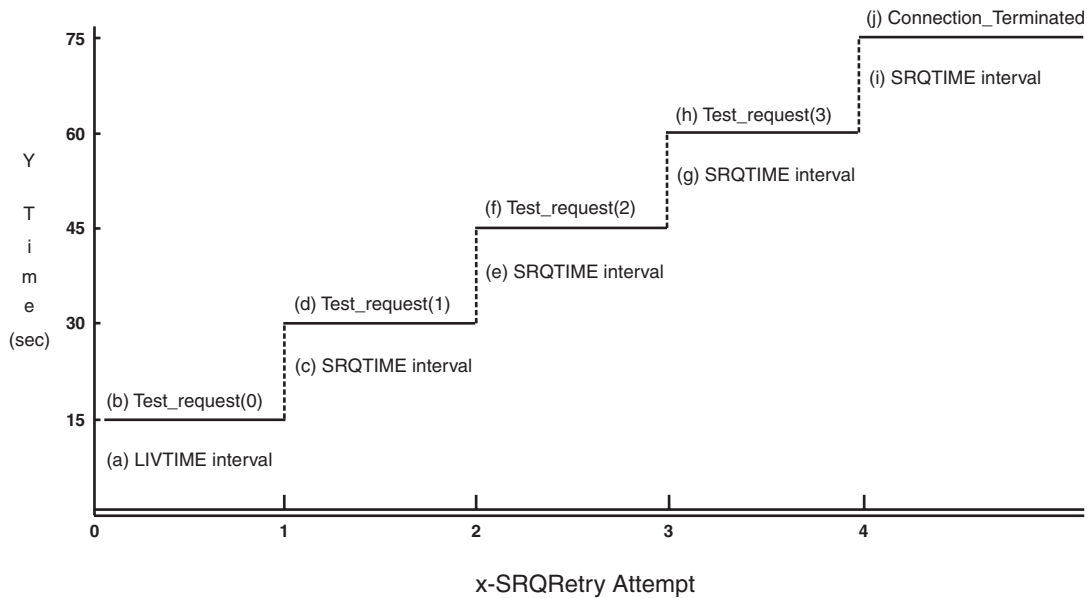


Figure 17. Enterprise Extender inactivity flows

- a. Initial LIVTIME interval expires and no inbound HPR traffic is detected by this EE endpoint during this 15-second interval.
- b. The Logical Data Link Control (LDLC) layer sends a TEST request to the EE partner to determine whether the partner is still there. This is the initial TEST, which is not considered a retry. The LIST=DETAIL output for the D EEDIAG,SRQRETRY command will display the initial test as attempt number zero in message IST2074I.
- c. After each TEST request is sent, VTAM waits a period of time equal to the SRQTIME value (15 seconds in this example).
- d. No response is received from the partner within the SRQTIME interval. The LDLC layer sends another TEST. This is considered the second retry attempt up to a maximum of SRQRETRY value (three retries for this example).
- e. After each TEST request is sent, VTAM waits a period of time equal to the SRQTIME value (15 seconds in this example).
- f. No response is received from the partner within the SRQTIME interval. The LDLC layer sends another TEST request. This is considered the third retry attempt, which is the final retry in this example.
- g. After each TEST request is sent, VTAM waits a period of time equal to the SRQTIME value (15 seconds in this example).
- h. No response is received from the partner within the SRQTIME interval. The LDLC layer sends another TEST request. This is considered the third retry attempt, which is the final retry in this example.
- i. After the final TEST request is sent, VTAM waits a period of time equal to the SRQTIME value (15 seconds in this example). Because this was the final retry attempt, and no response was received from the partner EE node, the EE connection is terminated with message IST1430I.

Using the DISPLAY EEDIAG output to tune your Enterprise Extender timers

The following sample display uses the same EE timer values as in Figure 17.

```

D NET,EEDIAG,SRQRETRY=3,LIST=DETAIL,CLEAR=SRQRETRY
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2066I ENTERPRISE EXTENDER CONNECTION SRQRETRY INFORMATION
.
.
IST2074I SUCCESSFUL SRQRETRY ATTEMPT = 0      OCCURRENCES =      98
IST2074I SUCCESSFUL SRQRETRY ATTEMPT = 1      OCCURRENCES =       5
IST2074I SUCCESSFUL SRQRETRY ATTEMPT = 2      OCCURRENCES =       2
IST2074I SUCCESSFUL SRQRETRY ATTEMPT = 3      OCCURRENCES =       1
.
.
IST314I END

```

The first IST2074I message displays attempt number 0 with 98 occurrences. To relate this to Figure 17 on page 134, this means the LIVTIME interval (a) expired without receiving any inbound HPR data, and a TEST request (b) was sent. The partner EE node responded to this TEST within the SRQTIME interval (c). Receiving a response is considered a successful attempt. Because this TEST is not considered a retry, it is denoted by SUCCESSFUL SRQRETRY ATTEMPT = 0. This scenario increments the number of occurrences for attempt number 0 by one. This display indicates that 98 times, when the initial TEST was sent out to the EE partner, VTAM received a reply from the partner within the first SRQTIME interval.

The second IST2074I message displays the attempt number of 1 with 5 occurrences. To relate this to Figure 17 on page 134, this means that the LIVTIME interval (a) expired without receiving any inbound HPR data, and a TEST request (b) was sent. No response was received from the partner EE node within the first SRQTIME interval (c). VTAM then sends out the first SRQRETRY attempt (d) for the TEST request. The partner EE node responded to this TEST within the SRQTIME interval (e). Receiving a response is considered a successful attempt. Because this TEST is considered a retry, it is denoted by SUCCESSFUL SRQRETRY ATTEMPT = 1. This scenario increments the number of occurrences for attempt number by one. This display indicates that five times when the first SRQRETRY attempt for the TEST was sent out to the EE partner, VTAM received a reply from the partner within this SRQTIME interval.

The remaining IST2074I messages indicate the same information, but the attempt number and number of occurrences indicates on which SRQRETRY VTAM received a response. Generally, the number of nonzero occurrences in the high numbered SRQRETRY attempts should be minimal, if not zero. If there are any SRQRETRY attempts in the last or next-to-last retry, this indicates that there are conditions in your IP network that are causing long network delays. In these situations, VTAM is very close to an inoperative condition on these EE connections.

The defaults for the LDLC parameters are probably sufficient for most networks, but tuning the parameters might be appropriate depending on the design of the underlying IP network and the technologies being used there. For example, if RIP is being used as the dynamic routing update protocol, then longer convergence times are to be expected (as compared to OSPF), and therefore the LDLC parameters could be adjusted (for example, by bumping the number of SRQRETRY attempts, or increasing the SRQTIME interval value to lengthen the time EE LDLC waits before inoping the connections).

Tuning your HPRPST values for your EE network

If an EE connection is suffering connectivity problems and is in the process of timing out (see “How Enterprise Extender times out an inactive connection” on page 133), the VTAM topology still thinks that EE is a viable route and might select it as the best route to that partner. If the EE link is truly experiencing problems (no inbound data in this case), RTP pipes eventually suffer problems as well. The RTP pipes will go into a path switch state fairly quickly. HPR path switch timers should be set long enough for the APPN topology to be updated to reflect the fact that the EE connection is no longer usable. This means that the HPRPST must be set longer than the time it takes for the EE connection to timeout. In Figure 17 on page 134, all the HPRPST values should be coded larger than 75S (75 seconds) to outlast the time it takes for an EE connection to time out because of inactivity.

Alternatively, you can specify HPRPSDLY=EEDELAY on the appropriate major node for your EE configuration. The HPRPSDLY parameter is available on the PU definition statement in the switched and model (DYNTYPE=EE) major nodes, and also on the connection network GROUP definition statements in the EE XCA major node. For more information about the HPRPSDLY parameter, see z/OS Communications Server: SNA Resource Definition Reference.

DISPLAY EEDIAG,TEST=YES

The DISPLAY EEDIAG,TEST=YES command, or Enterprise Extender connectivity test command, is useful for debugging network problems. Use this command to test an existing Enterprise Extender connection, or to assist in diagnosing why an EE connection cannot be established.

The EE connectivity test verifies EE line availability, address resolution capability, and ultimately partner reachability. Specify the DISPLAY EEDIAG,TEST=YES,LIST=DETAIL command to validate partner reachability. UDP requests with varying TTL (time-to-live) or hop count values are sent to the EE partner host. The command then waits for the routers between the local and remote hosts to send ICMP messages that indicate that the TTL value has been exceeded. If these messages are not received, the command provides the maximum number of retry attempts for that particular hop in the route. The DISPLAY EEDIAG,TEST=YES,LIST=SUMMARY connectivity test makes up to three attempts to reach the remote partner after VTAM sets the TTL count to 255. Because VTAM sets the TTL count to 255, the hop count is not determined for the LIST=SUMMARY output. Message IST2137I or IST2138I displays the hop count of *NA.

The output generated for the DISPLAY EEDIAG,TEST=YES,LIST=SUMMARY command lists the remote partner reachability information quickly.

The output generated from this request shows the reachability of the remote EE endpoint over all five UDP ports reserved for EE. When multipath routing or policy-based routing is being used, all available routes to the remote EE endpoint that are calculated by the local TCP/IP stack are tested.

New EE connection will not activate

Firewalls between the Enterprise Extender nodes must permit UDP traffic on all five EE ports for the IP address associated with each EE endpoint. If they do not, Enterprise Extender is not able to communicate. Figure 18 on page 137 depicts a

simple configuration that shows that the firewall protecting HostA is correctly configured to allow UDP traffic to pass through. The firewall guarding HostB has incorrectly left the firewall blocking UDP traffic. In this example, the firewalls are configured so that ICMP messages can pass through.

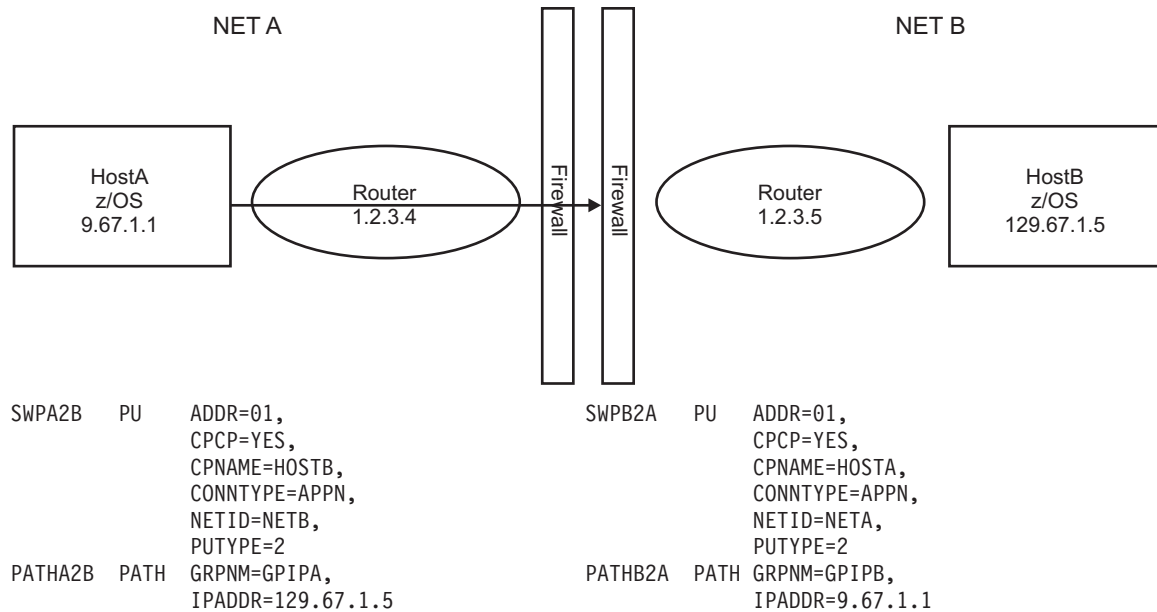


Figure 18. Enterprise Extender configuration with firewalls

You can use the EE connectivity test to assist you when a new EE connection will not activate. In Figure 18, the operator on HostA tries to dial a predefined switched PU to establish an EE connection to HostB. After the initial VARY ACCEPTED message, there is a pause, then an INOP occurs on the EE connection with reason XID OR LDLC COMMAND TIMEOUT. The INOP occurs in this case because the XIDs sent from HostA are not receiving responses back from HostB; the HostB firewall is discarding the XID packets. HostA will attempt to contact the partner up to the number of times specified by SRQRETRY. If no response is received for any XID, the connection fails.

Sample console log of the dial failure:

```
V NET,DIAL,ID=SWPA2B
```

```

IST097I VARY ACCEPTED
.
.
IST1411I INOP GENERATED FOR LNIP1
IST1430I REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT
IST314I END
  
```

Analyzing the problem

There are several reasons why the EE connection will not activate. The IPADDR or HOSTNAME value coded on the switched PU's PATH statement might not be coded correctly. There might be connectivity issues within the IP network that do not allow the EE connection to be established. If you have reviewed your Enterprise Extender definitions for accuracy, and you think that everything you have coded is correct, the problem might be within the IP network.

Performing an EE connectivity test

In the example shown in Figure 19, the operator on host HostA performs an EE connectivity test to assist in diagnosing the problem. In this example, there is no policy-based routing rule that matches the EE connection that is attempting to be established.

```
D NET,EEDIAG,TEST=YES,ID=SWPA2B,LIST=DETAIL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2067I EEDIAG DISPLAY ISSUED ON 08/21/05 AT 21:07:01
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST2023I CONNECTED TO LINE LN11
IST2126I CONNECTIVITY TEST IN PROGRESS
IST314I END

IST350I DISPLAY TYPE = EEDIAG
IST2130I ENTERPRISE EXTENDER CONNECTIVITY TEST INFORMATION
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2131I EEDIAG DISPLAY COMPLETED ON 08/21/05 AT 21:08:02
IST2132I LDLC PROBE VERSIONS: VTAM = V1          PARTNER = UNKNOWN
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST924I -----
IST2133I INTFNAME: LTRLE1A          INTFTYPE: MPCPTP
IST2135I CONNECTIVITY UNSUCCESSFUL  SENSE: ***NA***  PORT: 12000
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4          RTT: 3
IST2137I 2 *          (3)  RTT: *N/A*
IST2137I 7 *          (2)  RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL  SENSE: ***NA***  PORT: 12001
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4          RTT: 4
IST2137I 2 *          (3)  RTT: *N/A*
IST2137I 7 *          (2)  RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL  SENSE: ***NA***  PORT: 12002
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4          RTT: 4
IST2137I 2 *          (3)  RTT: *N/A*
IST2137I 7 *          (2)  RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL  SENSE: ***NA***  PORT: 12003
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4          RTT: 4
IST2137I 2 *          (3)  RTT: *N/A*
IST2137I 7 *          (2)  RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL  SENSE: ***NA***  PORT: 12004
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4          RTT: 4
IST2137I 2 *          (3)  RTT: *N/A*
IST2137I 7 *          (2)  RTT: *N/A*
IST924I -----
IST2139I CONNECTIVITY TEST INFORMATION DISPLAYED FOR 1 OF 1 ROUTES
IST314I END
```

Figure 19. Connectivity test without policy-based routing enabled

The example shown in Figure 20 on page 139 is identical to the example in Figure 19 except that a single policy-based routing rule (EEROUTINGRULE1) is being used for all EE traffic between the endpoints being tested. The policy-based routing rule has indicated that there is a route table defined for EE traffic (the route table name is EETABLE1).

```

D NET,EEDIAG,TEST=YES,ID=SWPA2B,LIST=DETAIL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2067I EEDIAG DISPLAY ISSUED ON 03/13/05 AT 21:07:01
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST2023I CONNECTED TO LINE LN11
IST2126I CONNECTIVITY TEST IN PROGRESS
IST314I END

IST350I DISPLAY TYPE = EEDIAG
IST2130I ENTERPRISE EXTENDER CONNECTIVITY TEST INFORMATION
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2131I EEDIAG DISPLAY COMPLETED ON 03/13/05 AT 21:07:46
IST2132I LDLC PROBE VERSIONS: VTAM = V1 PARTNER = UNKNOWN
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST2224I ENTERPRISE EXTENDER ROUTING POLICY INFORMATION
IST2225I PORT    ROUTE TABLE  ROUTING RULE
IST2205I ----    -
IST2226I 12000  EETABLE1      EEROUTINGRULE1
IST2226I 12001  EETABLE1      EEROUTINGRULE1
IST2226I 12002  EETABLE1      EEROUTINGRULE1
IST2226I 12003  EETABLE1      EEROUTINGRULE1
IST2226I 12004  EETABLE1      EEROUTINGRULE1
IST924I -----
IST2133I INTFNAME: TRLE1A          INTFTYPE: MPCPTP
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12000
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4                RTT: 22
IST2137I 2                      * (3) RTT: *N/A*
IST2137I 6                      * (3) RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12001
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4                RTT: 29
IST2137I 2                      * (3) RTT: *N/A*
IST2137I 6                      * (3) RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12002
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4                RTT: 23
IST2137I 2                      * (3) RTT: *N/A*
IST2137I 6                      * (3) RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12003
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4                RTT: 32
IST2137I 2                      * (3) RTT: *N/A*
IST2137I 6                      * (3) RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12004
IST2136I CONNECTIVITY TEST ENDED - MAXIMUM TIME LIMIT EXCEEDED
IST2137I 1 1.2.3.4                RTT: 22
IST2137I 2                      * (3) RTT: *N/A*
IST2137I 6                      * (3) RTT: *N/A*
IST924I -----
IST2139I CONNECTIVITY TEST INFORMATION DISPLAYED FOR 1 OF 1 ROUTES
IST314I END

```

Figure 20. EE connectivity test with a single policy-based routing rule enabled

Performing a basic EE connectivity test

In the example shown in Figure 21 on page 140, the operator on host HostA performs a basic EE connectivity test to determine the connectivity to the remote

host quickly. This example shows successful connectivity over all EE ports.

```
d net,eediag,id=swpa2b,test=yes,list=summary
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2067I EEDIAG DISPLAY ISSUED ON 06/30/10 AT 10:16:59
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST2023I CONNECTED TO LINE LNIP8
IST2126I CONNECTIVITY TEST IN PROGRESS
IST314I END
IST350I DISPLAY TYPE = EEDIAG
IST2130I ENTERPRISE EXTENDER CONNECTIVITY TEST INFORMATION
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2131I EEDIAG DISPLAY COMPLETED ON 06/30/10 AT 10:17:00
IST2132I LDLC PROBE VERSIONS: VTAM = V1          PARTNER = V1
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST924I -----
IST2133I INTFNAME: LCTC400          INTFTYPE: CTC
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12000
IST2137I *NA 129.67.1.5          RTT: 1
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12001
IST2137I *NA 129.67.1.5          RTT: 1
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12002
IST2137I *NA 129.67.1.5          RTT: 1
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12003
IST2137I *NA 129.67.1.5          RTT: 1
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12004
IST2137I *NA 129.67.1.5          RTT: 1
IST924I -----
IST2139I CONNECTIVITY TEST RESULTS DISPLAYED FOR 1 OF 1 ROUTES
IST314I END
```

Figure 21. Basic EE connectivity test (successful connection)

In the example shown in Figure 22 on page 141, the operator on host HostA performs a basic EE connectivity test to determine the connectivity to the remote host quickly. This example shows unsuccessful connectivity over EE ports 12003 and 12004.


```

d net,eediag,id=swpa2b,test=yes,list=summary
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000002
IST2067I EEDIAG DISPLAY ISSUED ON 06/30/10 AT 10:31:24
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST2023I CONNECTED TO LINE LNIP8
IST2126I CONNECTIVITY TEST IN PROGRESS
IST314I END
IST350I DISPLAY TYPE = EEDIAG
IST2130I ENTERPRISE EXTENDER CONNECTIVITY TEST INFORMATION
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000002
IST2131I EEDIAG DISPLAY COMPLETED ON 06/30/10 AT 10:31:35
IST2132I LDLC PROBE VERSIONS: VTAM = V1          PARTNER = UNKNOWN
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 129.67.1.5
IST924I -----
IST2133I INTFNAME: LCTC400          INTFTYPE: CTC
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12000
IST2137I *NA 129.67.1.5          RTT: 1
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12001
IST2137I *NA 129.67.1.5          RTT: 1
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12002
IST2137I *NA 129.67.1.5          RTT: 1
IST2135I CONNECTIVITY UNSUCCESSFUL      SENSE: ***NA***  PORT: 12003
IST2137I *NA *          RTT: *N/A*
IST2135I CONNECTIVITY UNSUCCESSFUL      SENSE: ***NA***  PORT: 12004
IST2137I *NA *          RTT: *N/A*
IST924I -----
IST2139I CONNECTIVITY TEST RESULTS DISPLAYED FOR 1 OF 1 ROUTES
IST314I END

```

Figure 22. Basic EE connectivity test (unsuccessful connection)

Figure 22 shows that ports 12003 and 12004 are blocked because of a firewall or some other reason. Issue the command `D NET,EEDIAG,ID=puname,TEST=YES,LIST=DETAIL` to diagnose the problem further.

Understanding the EE connectivity test output

Because the EE connectivity test is a potentially long-running command, the display output is broken into multiple sections. Some of the key messages in the first message group follow:

- Message IST2119I displays a unique display correlator that can be used to coordinate the various message groups of the DISPLAY EEDIAG command.
- Message IST2067I displays the date and time when the DISPLAY EEDIAG command was issued.
- Message IST1680I displays the local and remote IP addresses of the EE connection that is being tested.
- Message IST2023I displays the EE line that was selected to perform the EE connectivity test.
- Message IST2126I is an informational message that indicates that the connectivity test has been initiated.

Some of the key messages in the second message group follow:

- Message IST2130I is the header message in the EE connectivity test information message group.
- Message IST2119I displays a unique display correlator that can be used to coordinate the various message groups of the DISPLAY EEDIAG command.

- Message IST2131I displays the date and time when the EE connectivity test completed.
- Message IST2132I displays the version of the LDLC probe that VTAM is using to perform the connectivity test. If the connectivity test is successful across at least one port, this message also contains the EE partner's LDLC probe version.
- Message IST1680I displays the local and remote IP addresses of the EE connection that is being tested.
- If a policy-based routing rule is defined for any EE traffic between the EE endpoints, then you will also receive the following messages:
 - Message IST2224I is a header message displayed when a policy-based routing rule applies to EE traffic between the EE endpoints that are being tested for connectivity.
 - Message IST2225I is a header for the display of EE UDP ports, route tables, and the policy routing rules when a policy-based routing rule applies to EE traffic.
 - Message IST2226I displays the EE UDP ports and their associated route tables and policy routing rules when a policy-based routing rule applies to EE traffic. If a policy-based routing rule is not defined for an EE UDP port, then the policy routing rule NONE is specified. When the main routing table is being used (either a policy routing rule does not exist or the routing action indicates that the main routing table is being used), then the EZBMAIN route table is specified.
- Message IST2133I displays the TCP/IP interface; when multipath routing or policy-based routing is being used, the EE connectivity test is performed over each TCP/IP interface that can be used to route EE traffic to the requested destination.
- Message IST2135I indicates that the EE connectivity test was unsuccessful over this specific EE port.
- Message IST2136I indicates that the EE connectivity test ended for this port because the limit specified by the MAXTIME value was exceeded.

The LIST=DETAIL connectivity test makes up to three attempts at contacting each hop in the route. The test for each specific hop (or TTL value) stops when a response is received from the hop. After a response is received, or after the third attempt, the TTL value is increased by 1 to test the next hop and the test continues. Message IST2137I displays the results of each hop test. To reduce redundant output, VTAM prints only the first hop that did not receive a response, and the last hop that did not receive a response.

The LIST=SUMMARY connectivity test makes up to three attempts to reach the remote partner after VTAM sets the TTL count to 255. Because VTAM sets the TTL count to 255, the hop count is not determined for the LIST=SUMMARY output. Message IST2137I displays the hop count of *NA.

The output in “Performing an EE connectivity test” on page 138 shows that the LDLC probe used to test the connection did not receive any responses after the TTL reached a value of 2. The TTL was incremented by a value of 1 and retested. This was repeated until the TTL reached a value of 6. At this time, the maximum time limit allowed for the EE connectivity test (MAXTIME) was exceeded and the test ended. See the DISPLAY EEDIAG command in z/OS Communications Server: SNA Operation for more information.

Solving EE connectivity problems

In the previous EE connectivity test example, the EE connectivity test indicates that the EE traffic (UDP datagrams) cannot make it past the first hop in the route. The results are consistent for all five EE ports that were tested. At this point in the problem diagnosis, focus on the first hop in the EE route. Examine this hop for connectivity problems. Next, verify the routing tables for accuracy, check the logs for dropped packets, and verify that any firewall in the EE route allows UDP traffic for all five EE ports. If network address translation (NAT) is being used for Enterprise Extender connections, verify that the routers or nodes performing the NAT functions are translating the IP addresses to the correct addresses.

RTP performance problems over EE with multipath routing enabled

If multipath routing is enabled on the TCP/IP stack, and multiple equal-cost routes exist to the partner EE node, then TCP/IP sends batches of EE packets across each of these routes using a round-robin schedule. If one of these routes cannot reach the partner EE node, then EE might not activate, or if it does, there is likely to be significant performance impacts.




Figure 23. Enterprise Extender with multipath routing enabled

The following configuration information applies to Figure 23.

- Multipath routing is enabled in HostA
- Two QDIO OSA interfaces are defined and active in HostA
- Two static and equally weighted IP routes to destination HostB are defined in HostA
 - The IP route through IP Cloud1 has a router in the path; the router has incorrect routing definitions to HostB. A router in IP Cloud1 with IP address 9.67.1.21 is the router that is unable to route to 9.67.1.6 (HostB).
 - The IP route through IP Cloud2 has correct routing definitions to 9.67.1.6 (HostB).

In Figure 23 on page 143, an RTP pipe is successfully established from HostA over EE to HostB. However, the sessions using this RTP pipe are experiencing poor performance. The operator issues a `D NET,EEDIAG,REXMIT` command, which reveals that this EE connection is experiencing a high percentage of retransmissions. In this example, when the HPR traffic is routed over the path that uses the router with IP address 9.67.1.3, the HPR packet is correctly routed to HostB. When the HPR packets are transmitted over the route that uses the router with IP address 9.67.1.2, the packet is incorrectly routed and is subsequently discarded. The high percentage of lost packets causes the RTP endpoints to report lost packets, which causes subsequent retransmissions. Excessive retransmissions significantly degrade HPR throughput, and can lead to HPR path switches, or in some cases HPR connection deactivation.

To disable multipath for EE without affecting other IP applications, code the VTAM start option `MULTIPATH=NO` or allow it to default. This will disable the multipath function in the stack for EE connections only. The multipath behavior for other IP applications will remain unchanged.

Using the EE connectivity test to verify multipath routing

In "EE connectivity test with multipath routing enabled", the operator on HostA verifies the Enterprise Extender multipath routing environment by performing the following EE connectivity test:

```
D NET,EEDIAG,TEST=YES,IPADDR=(9.67.1.1,9.67.1.6),LIST=DETAIL
```

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE00000E
IST2067I EEDIAG DISPLAY ISSUED ON 10/04/05 AT 11:05:50
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.6
IST2023I CONNECTED TO LINE LN11
IST2126I CONNECTIVITY TEST IN PROGRESS
IST314I END
```

Figure 24. EE connectivity test with multipath routing enabled (part 1 of 2)

```

IST350I DISPLAY TYPE = EEDIAG
IST2130I ENTERPRISE EXTENDER CONNECTIVITY TEST INFORMATION
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE00000E
IST2131I EEDIAG DISPLAY COMPLETED ON 10/04/05 AT 11:05:52
IST2132I LDLC PROBE VERSIONS: VTAM = V1 PARTNER = V1
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.6
IST924I -----
IST2133I INTFNAME: OSA1          INTFTYPE: OSAFDDI
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12000
IST2137I 1 9.67.1.2              RTT: 2
IST2137I 2 9.67.1.21            D-1 RTT: 3
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12001
IST2137I 1 9.67.1.2              RTT: 2
IST2137I 2 9.67.1.21            D-1 RTT: 3
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12002
IST2137I 1 9.67.1.2              RTT: 2
IST2137I 2 9.67.1.21            D-1 RTT: 4
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12003
IST2137I 1 9.67.1.2              RTT: 2
IST2137I 2 9.67.1.21            D-1 RTT: 4
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: ***NA*** PORT: 12004
IST2137I 1 9.67.1.2              RTT: 2
IST2137I 2 9.67.1.21            D-1 RTT: 3
IST924I -----
IST2133I INTFNAME: OSA2          INTFTYPE: OSAFDDI
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12000
IST2137I 1 9.67.1.3              RTT: 9
IST2137I 2 9.67.1.11            RTT: 14
IST2137I 3 9.67.1.12            RTT: 19
IST2137I 4 9.67.1.4              RTT: 23
IST2137I 5 9.67.1.6              RTT: 27
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12001
IST2137I 1 9.67.1.3              RTT: 8
IST2137I 2 9.67.1.11            RTT: 14
IST2137I 3 9.67.1.12            RTT: 17
IST2137I 4 9.67.1.5              RTT: 21
IST2137I 5 9.67.1.6              RTT: 25
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12002
IST2137I 1 9.67.1.3              RTT: 8
IST2137I 2 9.67.1.11            RTT: 13
IST2137I 3 9.67.1.12            RTT: 18
IST2137I 4 9.67.1.5              RTT: 22
IST2137I 5 9.67.1.6              RTT: 27
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12003
IST2137I 1 9.67.1.3              RTT: 9
IST2137I 2 9.67.1.11            RTT: 19
IST2137I 3 9.67.1.12            RTT: 22
IST2137I 4 9.67.1.4              RTT: 24
IST2137I 5 9.67.1.6              RTT: 27
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12004
IST2137I 1 9.67.1.3              RTT: 7
IST2137I 2 9.67.1.11            RTT: 11
IST2137I 3 9.67.1.12            RTT: 12
IST2137I 4 9.67.1.4              RTT: 17
IST2137I 5 9.67.1.6              RTT: 23
IST924I -----
IST2139I CONNECTIVITY TEST INFORMATION DISPLAYED FOR 2 OF 2 ROUTES
IST314I END

```

Figure 25. EE connectivity test with multipath routing enabled (part 2 of 2)

This example clearly shows that connectivity from HostA to HostB over the OSA1 interface does not exist. The router with IP address 9.67.1.21 is returning an ICMP

message to the LDLC probe, which indicates the destination host is unreachable. For this case, investigate this router to determine why it returned this type of ICMP message.

The output also indicates that connectivity from HostA to HostB over the OSA2 interface does exist. Message IST2137I indicates that the route is a 5-hop configuration to the partner host. The display also shows that there are different routes through the IP network to the EE partner. When routing over the OSA2 interface, all five EE ports have successfully contacted the partner HostB with excellent round-trip times (RTT.)

EE connection or RTP pipe fails to activate when using policy-based routing for EE traffic

When policy-based routing is defined with multiple policy routing rules to separate traffic, then multiple routes can be used between the EE endpoints (even when multipath routing is not being used). If one of these routes cannot reach the partner EE node, then the EE connection might not activate; if it does activate, then one or more RTP pipes might not activate.

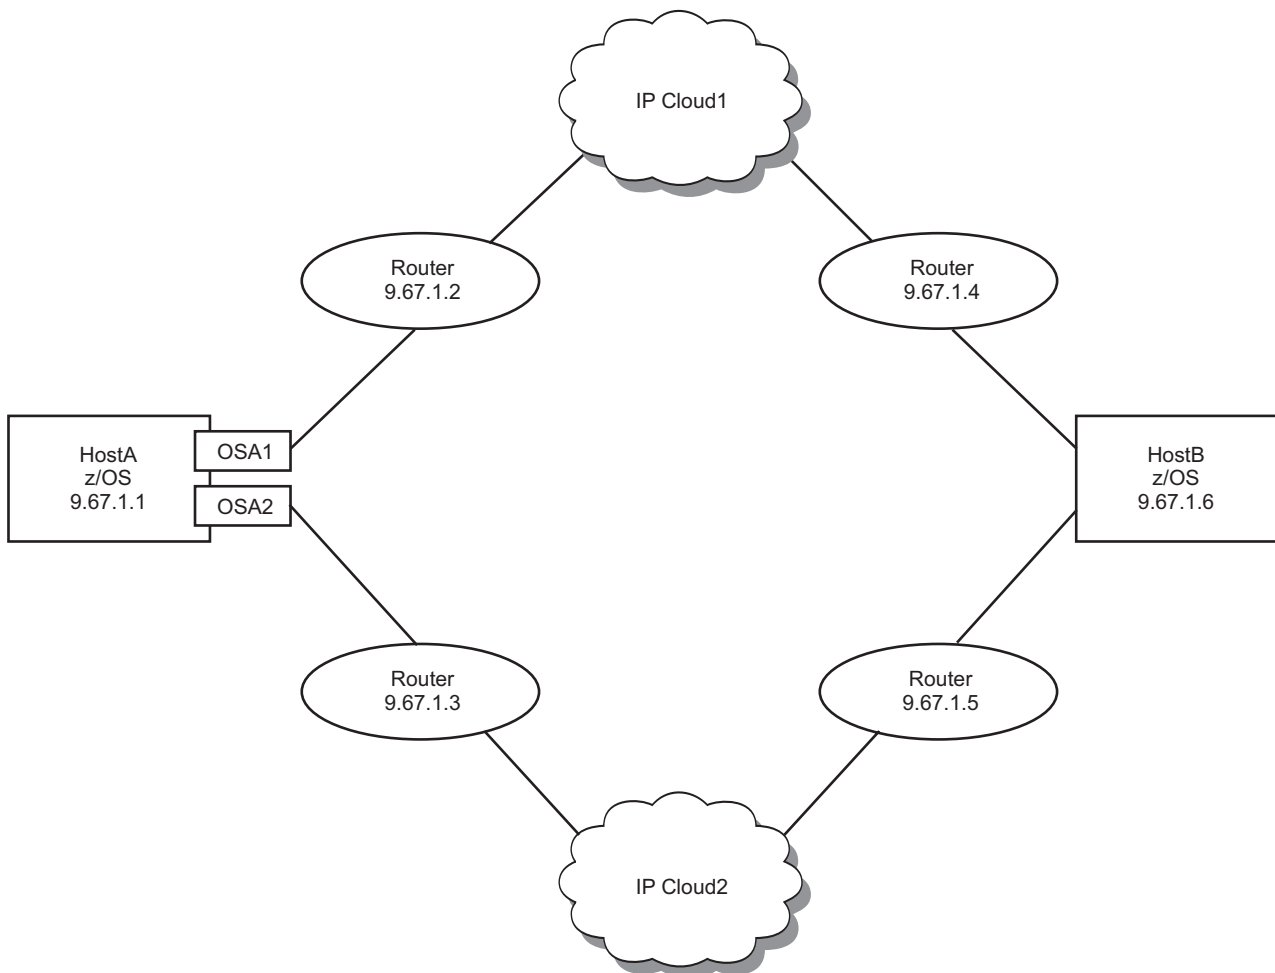


Figure 26. Enterprise Extender with policy-based routing

The following configuration information applies to Figure 26:

- Policy-based routing is being used with the following definitions:
 - A routing rule with the name EEROUTINGRULE1 is defined for EE traffic using EE UDP ports 12003 and 12004 (medium and low traffic priority data). This routing rule is associated with a routing action that points to route table EETABLE1. This route table has a statically defined IP route that uses the OSA1 interface and transmits all data to a next-hop IPv4 address 9.67.1.2 (into IP Cloud1).
 - A routing rule with the name EEROUTINGRULE2 is defined for EE traffic that uses EE UDP ports 12000, 12001, and 12002 (LDLC signal; network and high traffic priority data). This routing rule is associated with a routing action that points to route table EETABLE2. This route table has a statically defined IP route using the OSA2 interface and transmits all data to a next-hop IPv4 address 9.67.1.3 (into IP Cloud2).
- Multipath routing is disabled in host HostA.
- Two QDIO OSA interfaces are defined and active in host HostA.
 - Two static routes are defined between host HostA and host HostB.
 - The IP route through IP Cloud1 has a router in the path that has incorrect routing definitions to host HostB. A router in IP Cloud1 with IP address 9.67.1.21 is the router that is unable to route to IP address 9.67.1.6 (HostB).
 - The IP route through IP Cloud2 has correct routing definitions to IP address 9.67.1.6 (HostB).

In this example, an EE connection is successfully established from host HostA to host HostB. RTP pipes can be established using a transmission priority of high or network (CP-CP and RSETUP RTP pipes can be established). However, RTP pipes for low and medium transmission priorities fail to establish. When HPR traffic is routed over the path that uses the router with IP address 9.67.1.3, the HPR packet is correctly routed to host HostB. When HPR traffic is routed over the path that uses the router with IP address 9.67.1.2, the packet is incorrectly routed and is subsequently discarded. Therefore, a user can never establish an RTP pipe for low and medium transmission priorities.

Using the EE connectivity test to verify policy-based routing

In Figure 27 on page 148, the operator on host HostA verifies the EE policy-based routing environment by performing the following EE connectivity test:

```

D NET,EEDIAG,TEST=YES,IPADDR=(9.67.1.1,9.67.1.6),LIST=DETAIL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE00000E
IST2067I EEDIAG DISPLAY ISSUED ON 04/04/05 AT 11:05:50
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.6
IST2023I CONNECTED TO LINE LN11
IST2126I CONNECTIVITY TEST IN PROGRESS
IST314I END

IST350I DISPLAY TYPE = EEDIAG
IST2130I ENTERPRISE EXTENDER CONNECTIVITY TEST INFORMATION
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE00000E
IST2131I EEDIAG DISPLAY COMPLETED ON 04/04/05 AT 11:05:52
IST2132I LDLC PROBE VERSIONS: VTAM = V1 PARTNER = V1
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.6
IST2224I ENTERPRISE EXTENDER ROUTING POLICY INFORMATION
IST2225I PORT    ROUTE TABLE  ROUTING RULE
IST2205I ----    -
IST2226I 12000    EETABLE2      EEROUTINGRULE2
IST2226I 12001    EETABLE2      EEROUTINGRULE2
IST2226I 12002    EETABLE2      EEROUTINGRULE2
IST2226I 12003    EETABLE1      EEROUTINGRULE1
IST2226I 12004    EETABLE1      EEROUTINGRULE1
IST924I -----
IST2133I INTFNAME: OSA1          INTFTYPE: OSAFDDI
IST2227I CONNECTIVITY NOT TESTED DUE TO ROUTING POLICY PORT: 12000
IST2227I CONNECTIVITY NOT TESTED DUE TO ROUTING POLICY PORT: 12001
IST2227I CONNECTIVITY NOT TESTED DUE TO ROUTING POLICY PORT: 12002
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: 00000000 PORT: 12003
IST2137I 1 9.67.1.2              RTT: 10
IST2137I 2 9.67.1.21             D-1 RTT: 18
IST2135I CONNECTIVITY UNSUCCESSFUL SENSE: 00000000 PORT: 12004
IST2137I 1 9.67.1.2              RTT: 11
IST2137I 2 9.67.1.21             D-1 RTT: 18
IST924I -----
IST2133I INTFNAME: OSA2          INTFTYPE: OSAFDDI
IST2134I CONNECTIVITY SUCCESSFUL PORT: 12000
IST2137I 1 9.67.1.3              RTT: 9
IST2137I 2 9.67.1.11             RTT: 14
IST2137I 3 9.67.1.12             RTT: 19
IST2137I 4 9.67.1.5              RTT: 21
IST2137I 5 9.67.1.6              RTT: 27
IST2134I CONNECTIVITY SUCCESSFUL PORT: 12001
IST2137I 1 9.67.1.3              RTT: 8
IST2137I 2 9.67.1.11             RTT: 14
IST2137I 3 9.67.1.12             RTT: 17
IST2137I 4 9.67.1.5              RTT: 21
IST2137I 5 9.67.1.6              RTT: 25
IST2134I CONNECTIVITY SUCCESSFUL PORT: 12002
IST2137I 1 9.67.1.3              RTT: 8
IST2137I 2 9.67.1.11             RTT: 13
IST2137I 3 9.67.1.12             RTT: 18
IST2137I 4 9.67.1.5              RTT: 22
IST2137I 5 9.67.1.6              RTT: 27
IST2227I CONNECTIVITY NOT TESTED DUE TO ROUTING POLICY PORT: 12003
IST2227I CONNECTIVITY NOT TESTED DUE TO ROUTING POLICY PORT: 12004
IST924I -----
IST2139I CONNECTIVITY TEST INFORMATION DISPLAYED FOR 2 OF 2 ROUTES
IST314I END

```

Figure 27. EE connectivity test with multiple policy-based routing rules enabled

The example shown in Figure 27 indicates that policy-based routing is being used (as indicated by messages IST2224I, IST2225I, IST2226I, and IST2227I). The

displayed output indicates that EE traffic was defined to be routed over the two OSA interfaces by the EE UDP port that is being used for data transmission. Message IST2227I indicates when a test is not performed for a specific route as a result of the policy-based routing definitions. In the example, all traffic routed over EE UDP ports 12003 and 12004 is routed through the OSA1 interface, and all traffic routed over EE UDP ports 12000, 12001, and 12002 is routed through the OSA2 interface.

DISPLAY EEDIAG,TEST=PENDING

The DISPLAY EEDIAG,TEST=PENDING command lists outstanding Enterprise Extender displays that are pending host name resolution (DISPLAY EE or DISPLAY EEDIAG) or pending EE connectivity test results. Both host name resolution and the EE connectivity tests are potentially long-running functions. Use the D EEDIAG command format to query the outstanding displays to obtain the status. Each pending display has a unique EE display correlator assigned when the display command was issued. Message IST2119I identifies this correlator value, which can be used to identify the outstanding display request. A description of some of the key messages in the display output follow:

- Message IST2145I is the first message in the message group of the pending EE display commands output.
- Message IST2067I displays the date and time when the DISPLAY EEDIAG command was issued.
- Message IST2147I displays a unique display correlator that can be used to coordinate the various message groups of the DISPLAY EEDIAG command. When the original DISPLAY EE or DISPLAY EEDIAG command was issued, the correlator was assigned in message IST2119I. Message IST2147I also displays the current state of command processing. For EE connectivity tests, this message also displays the Enterprise Extender line being used to conduct this test.

Display Enterprise Extender connection network unreachable partner information

You can use the DISPLAY TOPO,LIST=UNRCHTIM command on a network node to obtain Enterprise Extender connection network unreachable partner information. The following information is displayed for each Enterprise Extender virtual node that has unreachable partner information:

- The network-qualified name of the Enterprise Extender virtual node
- The total number of unreachable partner paths associated with the virtual node
- If the unreachable partner limit is exceeded for the virtual node, an indication that it is exceeded, along with the lower unreachable partner threshold that must be reached before the virtual node will again be used
- The network-qualified name of the origin node on the unreachable paths through the connection network
- The network-qualified name of the unreachable partner (destination) on the unreachable paths through the connection network
- The unreachable time value specified for the Enterprise Extender connection network
- The time the unreachable timer expires for the unreachable paths through the connection network

The following commands provide information about EE connection network unreachable partners:

- To display all Enterprise Extender connection network unreachable partner information, use the following command: `D NET,TOPO,LIST=UNRCHTIM`
- To display Enterprise Extender connection network unreachable partner information associated with a specific virtual node, use the following command: `D NET,TOPO,LIST=UNRCHTIM,VRN=cp_name`
- To display Enterprise Extender connection network unreachable partner information associated with a specific origin node, use the following command: `D NET,TOPO,LIST=UNRCHTIM,ORIG=cp_name`
- To display Enterprise Extender connection network unreachable partner information associated with a specific partner node, use the following command: `D NET,TOPO,LIST=UNRCHTIM,DEST=cp_name`

The ORIG, VRN, and DEST operands can be used in any combination to control the scope of the unreachable partner information that is displayed. Depending on the value of the DSPLYWLD start option, you can use wildcard values for the ORIG, VRN, and DEST operands.

Display HPR route test

You can use the DISPLAY RTPS command with the ID or TCID operand and the TEST operand to test the performance characteristics of an RTP connection that has an endpoint in this VTAM node. When an HPR route test is performed, the results are displayed asynchronously at the console. These results show how long it took a route test packet (a diagnostic type of data packet) to traverse each link in the RTP path, as well as how long it took such a packet to travel from this end of the RTP connection to the other end. Thus, you can identify if any links are congested. A sample sequence of how you might go about diagnosing such a problem is shown:

- To display all RTP connections with an endpoint in this VTAM, issue the command:

```
D NET,RTPS
```

- Message group IST1695I is displayed, containing one instance of message IST1960I for each RTP connection with an endpoint in this VTAM. IST1960I gives the following information about each connection:

PU NAME

Name of the RTP PU used in this VTAM host as the ALS for this RTP connection

CP NAME

CP name of the host at the other end of the RTP connection

COS NAME

Class of Service (CoS) name for the sessions using this connection

SWITCH

Indicates whether a path switch is in progress (YES or NO)

CONGEST

Indicates that the connection is congested (YES or NO)

STALL

Indicates that the connection is stalled (YES or NO)

SESS

Number of sessions using this connection

- Issue the DISPLAY RTP command again with the TEST=YES operand and specify a particular RTP connection (by PU name or local TCID) to request an HPR route test for that RTP connection:

```
D NET,RTPS,ID=puname,TEST=YES
```

or

```
D NET,RTPS,TCID=local_tcid,TEST=YES
```

Tip: The TCID operand can be used to correlate a local RTP PU name to the RTP PU name used by the remote partner RTP node to represent the same RTP connection. To determine the RTP PU name used by the remote partner RTP node, first issue the DISPLAY NET,ID=*puname* command on the local node and remember the REMOTE TCID value shown on the end of message IST1476I. Then from the remote partner RTP node (shown on the IST1481I message of the prior display), issue the DISPLAY RTPS,TCID=*tcid* command using the REMOTE TCID value obtained from the prior display. The TEST=YES operand can also be included on this command.

- Message group IST1695I is again displayed, but this time with only one instance of message IST1960I that describes the identified RTP connection. Additionally, message IST1786I is issued confirming that an HPR route test is being initiated.
- When the HPR route test completes, message group IST1787I is displayed. It contains an instance of message IST1790I for each hop (link) in the RTP connection. IST1790I contains the following information for the link:

CP NAME

CP name of the node on the near side of the link

TG NUMBER

Transmission group number

PARTNER CP NAME

CP name on the far side of the link

INTERNODAL TIME

Time, in milliseconds, needed by a route test packet to traverse this link

- Also in message group IST1787I, message IST1792I provides the total time, in milliseconds, required for a route test packet to travel from this end of the RTP connection to the other end.

If a particular link or some of the links in the RTP connection appear to be slower than the others, you might want to take corrective action to alleviate the congestion problem on that link or links. For example, the following conditions might reveal a problem in routing at an intermediate node.

- The internodal time between an intermediate node in an RTP connection and the next node further from the origin is derived by subtracting the round-trip traversal time recorded for the packet sent to the intermediate node from the round-trip traversal time recorded for the packet sent to the next node. If the packet sent to the further node returned sooner, a minimum internodal time of 1 millisecond is set for the hop between the intermediate node and the further node, because it must be assumed that the packet did take a positive amount of time to travel from the intermediate node to the further node.
- The total RTP connection traversal time in message IST1792I is calculated by dividing by 2 the end-to-end traversal time recorded for the packet sent to the node at the other end of the RTP connection. In a case where this packet returned sooner than a packet to an intermediate node, the total RTP connection

traversal time shown in IST1792I will be less than the sum of the internodal times displayed in the IST1790I messages.

Display ID for an RTP connection

You can use the DISPLAY ID command to get information about an RTP connection. The resources that can be displayed and the most useful output messages are as follows:

- To display the RTP major node, use the command:

```
D NET, ID=ISTRTPMN
```

- Message IST1487I displays information about RTP ALS resources subordinate to the RTP major node:

RTP NAME

The RTP ALS name.

STATE

The connection state of the RTP ALS. Two states are presented: CONNECTED and CONNECTED/PSWITCH.

DESTINATION CP

The CPNAME of the adjacent RTP edge node.

TYPE The RTP connection type: LULU for RTPs with LU-LU sessions, RSTP for Route_Setup RTPs, and CPCP for RTPs with CP-CP sessions.

- To display an RTP ALS, use the command:

```
D NET, ID=rtp_als_name
```

- Message IST1479I displays the RTP connection state.
- Message IST1461I displays the portion of the session path managed by the RTP connection.
- Message IST875I displays the adjacent link station for the RTP connection.
- Messages IST1738I and IST1739I display the automatic network routing (ANR) labels and corresponding transmission priorities and explicit route numbers for the RTP connection.

- To display an ADJCP, use the command:

```
D NET, ID=adjcp_name, ADJCP
```

- Message IST1487I displays the RTP connections related to this ADJCP.

Note: A DISPLAY for an ADJCP representing a physically adjacent node also issues messages IST1106I presenting information on DLC-level PUs. A DISPLAY for a logically adjacent ADJCP, representing a distant RTP end-point, issues only messages IST1487I.

Display ID for an RTP PU with HPRDIAG=YES

You can use the DISPLAY ID command to get information about an RTP connection. To display the HPR diagnostic information for the RTP physical unit, specify the HPRDIAG=YES option.

Message IST2244I displays the date and time the DISPLAY ID command with HPRDIAG=YES was issued.

Several messages display the information about the RTP pipe.

ARB information:

- Message IST1844I displays the ARB mode.
- Messages IST1477I, IST1516I, IST1697I, IST1841I, IST1846I, IST1862I, IST2267I, and IST2395I display ARB information.
- Message IST1969I displays the maximum actual data flow rate since the last time counters were cleared.
- Message IST1970I displays the rate reductions because of retransmission since the last time counters were cleared.

Timer information:

- Messages IST1852I, IST1851I, IST1972I, and IST2229I display the timer information.

Outbound transmission information:

- Messages IST1974I, IST1975I, and IST1980I display the information about the outbound transmission since the last time the counters were cleared.
- Message IST1980I displays the sequence number of the last received byte.
- Message IST1842I displays the number of NLPS retransmitted since the last time the counters were cleared.
- Message IST2249I displays the NLP retransmit rate since the last time the counters were cleared.
- Message IST2236I displays the time the last NLP was retransmitted.
- Message IST1976I displays the number of bytes retransmitted since the last time the counters were cleared.
- Message IST1478I displays the number of unacknowledged buffers.
- Message IST1958I displays the number of orphaned buffers since the last time the counters were cleared.
- Messages IST1843I, IST1847I, IST2085I, and IST1511I display additional information about the outbound transmission.
- Messages IST1977I, IST1978I, IST2086I, and IST2087I display additional information about the outbound transmission since the last time the counters were cleared.

Inbound transmission information:

- Message IST2059I displays the number of NLPS received since the last time the counters were cleared.
- Message IST1981I displays the total number of bytes received since the last time the counters were cleared.
- Message IST1850I displays the largest NLP received since the last time the counters were cleared.
- Message IST2230I displays the maximum number of NLPS on the out of sequence queue since the last time it was reset to the current number of NLPS on the out of sequence queue.
- Messages IST1980I, IST1853I, IST1854I, and IST1982I display additional information about the inbound transmission.
- Message IST1983I displays the maximum number of NLPS on inbound work queue since the last time it was reset to the current number of NLPS on inbound work queue.

Path switch information:

- Messages IST1856I, IST1937I, IST1985I, IST1986I, IST1987I, and IST1988I display information about the path switch since the last time the counters were cleared. These messages are not displayed if there was no path switch since the last time the counters were cleared.

Back pressure reason counts:

- Messages IST1858I, IST1859I, IST2211I, IST2212I, IST2213I, and IST2215I display information about the back pressure since the last time the counters were cleared. These messages will not be displayed if there was no back pressure since the last time the counters were cleared.

Last time diagnostic counters cleared message:

- Message IST2250I displays the date and time when the diagnostic counters were last cleared.

Display ID for an RTP PU with HPRDIAG=YES and CLEAR=ALL

The following diagnostic counters are cleared or reset after the DISPLAY command output. They are displayed before being cleared by the message shown on DISPLAY ID of the RTP PU with HPRDIAG=YES command:

- The high water mark of the smooth sending rate is reset to the current smooth sending rate. It is displayed by the message IST1969I.
- The number of rate reductions due to retransmission is cleared. It is displayed by the message IST1970I.
- The count of NLPs sent is cleared. It is displayed by the message IST1974I.
- Total number of bytes sent is cleared. It is displayed by the message IST1975I.
- Largest NLP sent is cleared and it is displayed by the message IST1849I.
- Number of NLPs retransmitted is cleared and it is displayed by the message IST1842I.
- Number of retransmitted bytes is cleared and it is displayed by the message IST1976I.
- Number of orphaned buffers is cleared and it is displayed by the message IST1958I.
- The high water mark for the number of NLPs waiting on the acknowledgment queue is reset to the current number of NLPs waiting on the acknowledgment queue. It is displayed by the message IST1977I.
- TOD clock of high water mark for the number of NLPs waiting on acknowledgment queue is cleared. It is displayed by the message IST1978I.
- The high water mark for the number of NLPs on the outbound work queue is reset to the current number of NLPs on the outbound work queue. It is displayed by the message IST2086I.
- TOD clock of high water mark for the number of NLPs on the outbound work queue is cleared. It is displayed by the message IST2087I.
- The number of NLPs received is cleared and it is displayed by the message IST2059I.
- The total number of bytes received is cleared and it is displayed by the message IST1981I.
- Largest NLP received is cleared and it is displayed by the message IST1850I.

- The maximum number of NLPs on the out of sequence queue is reset to the number of NLPs on the in of sequence queue. It is displayed by the message IST2230I.
- The maximum number of NLPs on the inbound work queue is reset to the number of NLPs on the inbound work queue. It is displayed by the message IST1983I.
- Path switches initiated from remote RTP is cleared. It is displayed by the message IST1985I.
- Path switches initiated from local RTP is cleared. It is displayed by the message IST1986I.
- Path switches initiated due to local failure is cleared. It is displayed by the message IST1987I.
- Path switches initiated due to local PSRETRY is cleared. It is displayed by the message IST1988I.
- Back pressure pathswitch count, back pressure sendq max count, back pressure storage failure count and back pressure stall count are cleared. They are displayed by the message IST1859I.
- Back pressure waiting for the acknowledgment maximum counter is cleared. It is displayed by the message IST2212I.
- TOD of the last back pressure applied is cleared and it is displayed by the message IST2213I.
- Last back pressure reason is cleared. It is displayed by the message IST2215I.

Display NCP storage

You can display any selected portion of NCP storage by using a DISPLAY NCPSTOR command. Up to 256 bytes can be displayed for each command. For the 3720 or 3745 communication controller, you can also use the DISPLAY NCPSTOR command to display up to 256 bytes of an NCP dump or state vector stored on the disk.

VTAM issues message IST245I to display the storage contents.

Note:

1. The NCP storage to be displayed might change while it is being formatted for transmission.
2. You cannot display NCP storage at a data host, because a data host does not own any NCPs.

Display path tables

This is the place to start when route problems are detected. This display provides information about the route status known by this host VTAM.

You can use the DISPLAY PATHTAB command to display the status of explicit routes and their associated virtual routes for a local host. You can display information about all routes or you can limit the information using the ADJSUB and DESTSUB operands. The resulting display shows the host path table contents.

Display resource status

You can use the DISPLAY ID command to display status information about any major or minor node. For example, a request to display a physical unit or a logical unit indicates whether that node has been added by dynamic reconfiguration. DISPLAY ID also indicates whether a logical unit, a physical unit, or a link is supported by the Network Terminal Option (NTO).

Note: From a data host, you cannot display the status of either an NCP or the NCP subordinate resources, because a data host does not own any NCPs.

Two types of node status are displayed when you use the DISPLAY ID command:

- The STATUS field shows the *current state* of the node
- The DESIRED STATE field shows the *desired state*

The desired state is the condition that VTAM processing is attempting to establish for the node. Previously entered operator commands or recovery processing can establish the desired state for a node. When processing is completed, the desired state and the current state should be the same.

If SNA network interconnection is in use, the DISPLAY command shows:

- The network ID associated with a resource (if any)
- For a cross-network CDRSC, the real resource name
- For a cross-network CDRM, the real name of the gateway node through which the SSCP-SSCP session passes, and the network address as known in the requesting host's network

Display resources in a pending state

You can use the DISPLAY PENDING command to display information about resources in the domain that are in one of the following pending states:

- Transient state to or from a fully active state.
- State of “recovery pending” or “recovery in progress” for application programs suspended because of the failure or takeover of an application program enabled for persistence. You *must* enter the DISPLAY PENDING command on the system in which the application program resides.

The resource can be a major node, a minor node, an application program, a physical unit, or a logical unit.

Display route status

The DISPLAY ROUTE command shows the status and availability of virtual and explicit routes. VTAM displays the status of selected routes and, if TEST=YES is specified, does a route test on the routes selected in the DISPLAY command. (See “Display route test” on page 163) The VTAM operator can select the origin of the routes to be displayed or tested. The origin can be either a host processor or an NCP.

For a sample path problem, see “Example: Solving path problems” on page 157.

The display of status for the routes selected is formatted as shown in Figure 28 on page 157.


```

Displaying one explicit route to a destination subarea:
d net,route,destsub=01,netid=netc,origin=a03n43a,er=5
IST097I DISPLAY ACCEPTED
IST535I ROUTE DISPLAY 7 FROM SA 4 TO SA 1
IST808I ORIGIN PU = C0453LE DEST PU = C01NPU NETID = NETC
IST536I VR TP STATUS ER ADJSUB TGN STATUS CUR MIN MAX
IST537I 0 0 ACTIV 5 1 1 ACTIV3
IST537I 0 1 INACT 5 1 1 ACTIV3
IST537I 0 2 INACT 5 1 1 ACTIV3
IST314I END

```

Figure 28. Example of DISPLAY route status output

If you are using SNA network interconnection, the DISPLAY ROUTE command can be used to show the status and availability of adjacent VTAM networks. The resulting display is the same as shown in the previous example, except that message IST808I contains an additional field, NETID=*netid*, to show the ID of the adjacent network.

See z/OS Communications Server: SNA Messages for a complete description of the variable data contained in the messages that result from the DISPLAY ROUTE command.

It is possible for the test results for an explicit route to be lost before they are displayed. For example, if a node or a link along the explicit route fails between the time the test request flows outbound and the time the test results flow inbound, the results will be lost. If this occurs, reenter the DISPLAY ROUTE command for that explicit route.

Note: You can use the NetView session monitor to collect more information about routes. If all the required session monitors along the route are in session, you can test the entire route, from one session end to the other. See Table 48 on page 649 to determine what document has more information about the NetView session monitor.

Example: Solving path problems

This example takes you through a sample path problem and shows you how to diagnose and solve the problem. It includes the following topics:

- “Rules for routing”
- “Configuration and situation” on page 158
- “Analyzing the problem” on page 158
- “Finding the problem” on page 159
- “Steps for displaying routes” on page 159
- “Fixing the problem” on page 162
- “Dynamic path update” on page 162
- “Coding the path in the NCP” on page 162

Rules for routing

The rules that you need to keep in mind when working with routes are:

- Virtual routes must end in the subareas where the session end points reside.

- Virtual route numbers must be defined the same in both directions but only at each end point. A virtual route definition does not need to be specified in every node in the path.
- Explicit routes do not have to flow in both directions, but must be the same in one direction from endpoint subarea to endpoint subarea.
- Explicit and virtual route rules apply in each network, not across SNI network boundaries.
- Transmission group numbers on the VR that you are using must be the same in both directions, but only between two nodes, not along the entire path.

Configuration and situation

The following figure describes the configuration:

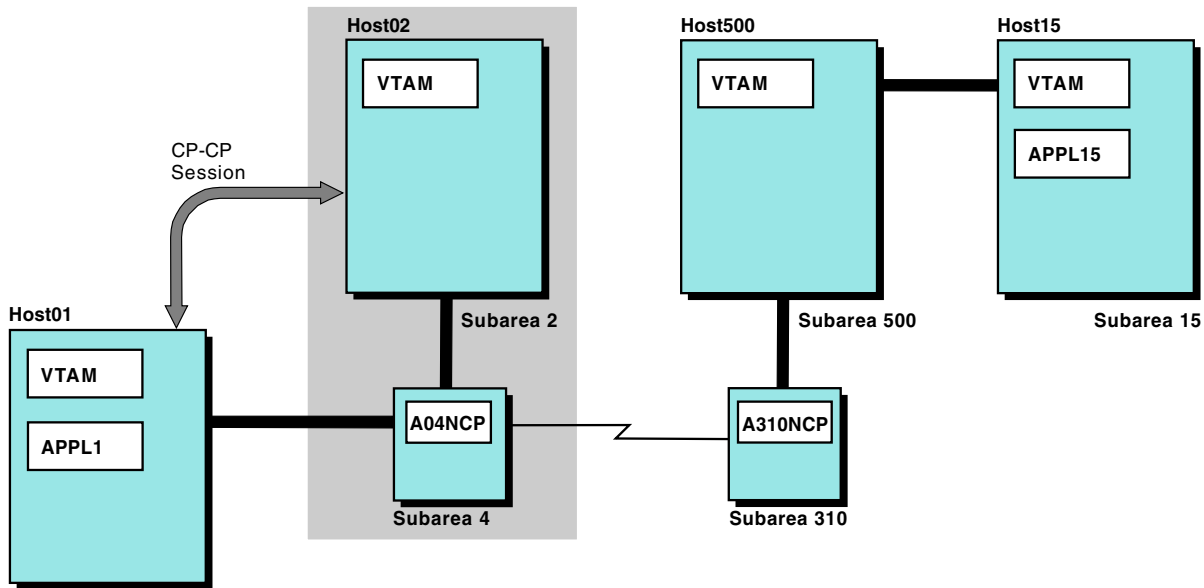


Figure 29. Path problem example network configuration

The situation is:

- Host01 is a network node and has a CP-CP session with Host02.
- Host15 and Host500 are subarea data hosts.
- You need a session from APPL1 to APPL15.
- The session is failing with a path problem sense code (8013xxxx).

Analyzing the problem

Because Host01 is an APPN node, the route does not end in Host01; it ends in subarea 4 (A04NCP). The route must follow the path: subarea 4 (A04NCP) to 310 (A310NCP) to 500 (Host500) to 15 (Host15) and back.

To follow the rules for routing, you need the following information for this session:

- One virtual route number going from Host500 to A04NCP
- One explicit route number going from Host15 to A04NCP
- One explicit route number going from A04NCP to Host15

Finding the problem

To find the problem, you must first display the routes across your session path. Then, you need to map the information that you received from the displays to locate the problem.

Steps for displaying routes

You can display the routes between the subareas in your routes to identify the problem area.

Note: Displays shown are abbreviated. If you run this display on your system, it will list information for all explicit routes from 0 to 15.

1. Display the route from Host15 to Host500 with A04NCP as the destination subarea. From Host15, issue the following command:

D NET,ROUTE,DESTSUB=4

```

IST097I DISPLAY ACCEPTED
IST535I ROUTE DISPLAY 1 FROM SA 15 TO SA 4 225
IST808I ORIGIN PU = ISTPUS DEST PU = A04NCP NETID = NETA
IST536I VR TP STATUS ER ADJSUB TGN STATUS CUR MIN MAX
IST537I 0 3 1 INOP
IST537I 1 0 ACTIV 1 500 2 ACTIV3 7 5 15
IST537I 1 1 INACT 1 500 2 ACTIV3
IST537I 1 2 ACTIV 1 500 2 ACTIV3 12 5 15
IST537I 3 0 INACT 2 500 1 INOP
IST537I 5 0 INACT 3 30 1 INOP
IST537I 5 1 INACT 3 30 1 INOP
IST537I 5 2 ACTIV 3 30 1 INOP
IST537I 6 0 INACT 6 500 1 ACTIV3
IST537I 6 1 INACT 6 500 1 ACTIV3
IST537I 6 2 ACTIV 6 500 1 ACTIV3 23 15 45
IST537I 15 UNDEF
IST314I END
  
```

From subarea 15 to subarea 500, you have ER1, ER2, and ER6 defined.



2. Display the route from Host500 to A310NCP with A04NCP as the destination subarea. From Host500, issue the following command:

D NET,ROUTE,DESTSUB=4

```


IST097I DISPLAY ACCEPTED
IST535I ROUTE DISPLAY 1 FROM SA 500 TO SA 4
IST808I ORIGIN PU = ISTPUS DEST PU = A04NCP NETID = NETA
IST536I VR TP STATUS ER ADJSUB TGN STATUS CUR MIN MAX
IST537I 0 2 1 INOP
IST537I 1 0 ACTIV 1 310 2 ACTIV3 20 15 45
IST537I 1 1 INACT 1 310 2 ACTIV3
IST537I 1 2 ACTIV 1 310 2 ACTIV3 33 15 45
IST537I 2 0 ACTIV 2 310 2 ACTIV3 9 5 15
IST537I 6 0 INACT 6 310 1 ACTIV1
IST537I 6 1 INACT 6 310 1 ACTIV1
  
```

```

IST537I 6 2 ACTIV 6 310 1 ACTIV3 7 5 15
IST537I 7 0 INACT 7 310 1 INOP
IST537I 15 UNDEF
IST314I END

```

From subarea 500 to subarea 310, you have ER1, ER2, ER6, and ER7 defined.



3. Display the route from A310NCP to A04NCP. From Host500, issue the following command:


D NET,ROUTE,ORIGIN=A310NCP,DESTSUB=04

```

IST097I DISPLAY ACCEPTED
IST535I ROUTE DISPLAY 1 FROM SA 310 TO SA 4 225
IST808I ORIGIN PU = A310NCP DEST PU = A04NCP NETID = NETA
IST536I VR TP STATUS ER ADJSUB TGN STATUS CUR MIN MAX
IST537I 0 2 1 INOP
IST537I 1 0 INACT 1 4 1 INOP
IST537I 6 0 INACT 6 4 1 ACTIV3
IST537I 6 1 INACT 6 4 1 ACTIV3
IST537I 6 2 ACTIV 6 4 1 ACTIV3 3 2 6
IST537I 15 UNDEF
IST314I END

```

From subarea 310 to subarea 4, you have ER1 and ER6 defined.



So, you can use either ER1 or ER6 to go from subarea 15 to subarea 4.

4. Display the route from A04NCP to A310NCP with Host15 as the destination subarea. From Host02, issue the following command:


D NET,ROUTE,ORIGIN=A04NCP,DESTSUB=15

```

IST097I DISPLAY ACCEPTED
IST535I ROUTE DISPLAY 1 FROM SA 4 TO SA 15 225
IST808I ORIGIN PU = A04NCP DEST PU = A15PU NETID = NETA
IST536I VR TP STATUS ER ADJSUB TGN STATUS CUR MIN MAX
IST537I 0 2 1 INOP
IST537I 2 UNDEF
IST537I 7 0 INACT 7 310 1 ACTIV3
IST537I 7 1 INACT 7 310 1 ACTIV3
IST537I 7 2 ACTIV 7 310 1 ACTIV3 3 2 6
IST537I 15 UNDEF
IST314I END

```

From subarea 4 to subarea 310, you have ER7 defined.



5. Display the route from A310NCP to Host500 with Host15 as the destination subarea. From Host500, issue the following command:


D NET,ROUTE,ORIGIN=A310NCP,DESTSUB=15

```

IST097I DISPLAY ACCEPTED
IST535I ROUTE DISPLAY 1 FROM SA 310 TO SA 15 225
IST808I ORIGIN PU = A310NCP DEST PU = A15PU NETID = NETA
IST536I VR TP STATUS ER ADJSUB TGN STATUS CUR MIN MAX
IST537I 0 2 1 INOP
IST537I 1 1 INACT 2 500 2 INOP
IST537I 2 0 INACT 1 500 1 ACTIV3 20 15 45
IST537I 7 0 ACTIV 7 500 1 ACTIV3 29 20 60
IST537I 7 1 ACTIV 7 500 1 ACTIV3
IST537I 7 2 ACTIV 7 500 1 ACTIV3 40 20 60
IST537I 15 UNDEF
IST314I END

```

From subarea 310 to subarea 500, you have ER1, ER2, and ER7 defined.



6. Display the route from Host500 to Host15. From Host500, issue the following command:


D NET,ROUTE,DESTSUB=15

```

IST097I DISPLAY ACCEPTED
IST535I ROUTE DISPLAY 1 FROM SA 500 TO SA 15 225
IST808I ORIGIN PU = ISTPUS DEST PU = A15PU NETID = NETA
IST536I VR TP STATUS ER ADJSUB TGN STATUS CUR MIN MAX
IST537I 0 2 1 INOP
IST537I 1 1 INACT 2 15 2 INOP
IST537I 4 0 INACT 5 310 1 INOP
IST537I 4 1 INACT 5 310 1 INOP
IST537I 4 2 INACT 5 310 1 INOP
IST537I 7 0 INACT 7 15 1 ACTIV3 7 5 15
IST537I 7 1 INACT 7 15 1 ACTIV3
IST537I 7 2 ACTIV 7 15 1 ACTIV3 35 20 60
IST537I 15 UNDEF
IST314I END


```

From subarea 500 to subarea 15, you have ER2 and ER7 defined.



Fixing the problem

Here is what your routing looks like now:



From these examples, you can see that the easiest way to fix the problem is to define ER2 from subarea 4 to subarea 310 and map VR1 to ER2. You can do this dynamically using dynamic path update, or you can change the NCP generation.

Dynamic path update

To fix this problem using dynamic path update, define ER2 and map VR1 to ER2.

Guideline: Whenever you change paths, make sure that you are not deleting a path that you need for another route. The following example is the NCPPATH statement that defines the new path.

```
A04NCP  NCPPATH NETID=NETA
P002    PATH DESTSA=310,ER2=(310,1),VR1=2
```

To change the path, use the VARY ACT command to activate your dynamic path update member.

Note: If this was a problem from a host to another subarea, you would use a VPATH definition to fix the problem

Coding the path in the NCP

In the NCP generation, find the path for destination subarea 310:

```
PATH DESTSA=310,
  ER0=(2,1),ER1=(71,80),
  ER3=(2,1),ER4=(3,80),ER5=(310,80),
  ER6=(310,80),ER7=(310,1),ER8=(71,80),
  ER9=(2,1),ER10=(1,1),ER11=(400,80),
  ER12=(1,1),
  VR0=6,
  VRPWS00=(1,3),VRPWS01=(1,3),VRPWS02=(1,3),
  VR1=9,
  VRPWS10=(2,6),VRPWS11=(2,6),VRPWS12=(2,6),
  VR2=3,
  VRPWS20=(2,6),VRPWS21=(2,6),VRPWS22=(2,6),
  VR3=8,
  VRPWS30=(2,6),VRPWS31=(2,6),VRPWS32=(2,6),
  VR4=4,
  VRPWS40=(2,6),VRPWS41=(2,6),VRPWS42=(2,6),
  VR5=11,
```

```

VRPWS50=(2,6),VRPWS51=(2,6),VRPWS52=(2,6),
VR6=10,
VRPWS60=(3,9),VRPWS61=(3,9),VRPWS62=(3,9),
VR7=7,
VRPWS70=(3,9),VRPWS71=(3,9),VRPWS72=(3,9)

```

Change the definition for ER2 to destination subarea 310 and TGN1 and map VR1 to ER2:

```

PATH DESTSA=310,
ER0=(2,1),ER1=(71,80),ER2=(310,1),
ER3=(2,1),ER4=(3,80),ER5=(310,80),
ER6=(310,80),ER7=(310,1),ER8=(71,80),
ER9=(2,1),ER10=(1,1),ER11=(400,80),
ER12=(1,1),
VR0=6,
VRPWS00=(1,3),VRPWS01=(1,3),VRPWS02=(1,3),
VR1=2
VRPWS10=(2,6),VRPWS11=(2,6),VRPWS12=(2,6),
VR2=3,
VRPWS20=(2,6),VRPWS21=(2,6),VRPWS22=(2,6),
VR3=8,
VRPWS30=(2,6),VRPWS31=(2,6),VRPWS32=(2,6),
VR4=4,
VRPWS40=(2,6),VRPWS41=(2,6),VRPWS42=(2,6),
VR5=11,
VRPWS50=(2,6),VRPWS51=(2,6),VRPWS52=(2,6),
VR6=10,
VRPWS60=(3,9),VRPWS61=(3,9),VRPWS62=(3,9),
VR7=7,
VRPWS70=(3,9),VRPWS71=(3,9),VRPWS72=(3,9)

```

Attention: Whenever you change paths, make sure that you are not deleting a path that you need for another route.

Display route test

If a route test was requested, results of the test are sent asynchronously to the console of the operator requesting the display. If the route test failed, the results are also sent to the console of the host that owns the rejecting subarea node. If the host owning the rejecting subarea is the same host that initiated the route test, that host will receive the test results twice.

To be tested, the explicit route must be known to VTAM. This means that the explicit route must be defined to VTAM, or at some time must have been operative.

Successful route test

If TEST=YES is set and the route test is successful, the following asynchronous messages follow the route status display messages previously described.

The test results are formatted as shown in Figure 30 on page 164.

```

IST538I  ROUTE TEST ### IN PROGRESS
IST533I  ER n SUCCEEDED IN ROUTE TEST ###
IST797I      FROM   VIA     ADJACENT   DEST   ER LENGTH
IST644I      ffffffff TG     aaaaaaaaa dddddddd
IST534I      sss    t       xxx       yyy    1
IST798I      nnnn

```

Figure 30. Output of a successful route test

In this example,

- The name of the origin physical unit is *fffffff*
- The adjacent node is *aaaaaaaa*.
- The name of the destination physical unit is *ddddddd*.
- The subarea number of *fffffff* is *sss*.
- The transmission group number is *t*.
- The subarea number of *aaaaaaaa* is *xxx*.
- The subarea of *ddddddd* is *yyy*.
- The explicit route length is 1.
- The network ID of the node being displayed is *nnnn*.

Failed route test

If the explicit route test fails because VTAM is unable to send the Explicit Route Test RU into the network, a message tells why the test cannot be performed. This message is shown in the following example.

```
IST510I  ROUTE TEST ### FAILED - reason
```

If the explicit route test is initiated by VTAM but fails, the messages in Figure 31 show the reason for the test failure.

```

IST533I  ER 0 FAILED IN ROUTE TEST 8
IST797I      FROM   VIA     ADJACENT   DEST   ER LENGTH
IST644I      ffffffff TG     aaaaaaaaa dddddddd
IST534I      sss    t       xxx       yyy    1
IST798I      nnnn
IST572I      REJECTING TG     ADJACENT   ER MASK
IST816I      rrr    g       zzz       mmmm
IST523I      <ER NOT DEFINED>
            <A REQUIRED TG IS INACTIVE>
            <ER NOT REVERSIBLE>
            <ER EXCEEDS MAXIMUM LENGTH>
            <MIGRATION ER NOT SUPPORTED>
            <MIGRATION NODE DOES NOT SUPPORT THIS ER>
            <MIGRATION NODE ENCOUNTERED>
            <UNEXPECTED TYPE BYTE X'##'>

```

Figure 31. Output of a failed route test

In this example,

- The name of the physical unit which originated the ER_TEST is *fffffff*.
- The adjacent physical unit is *aaaaaaaa*.
- The name of the destination physical unit is *ddddddd*.
- The subarea number of *fffffff* is *sss*.
- The transmission group number is *t*.
- The subarea number of *aaaaaaaa* is *xxx*.
- The subarea of *ddddddd* is *yyy*.

- The explicit route length is 1.
- The network ID of the node being displayed is *nnnn*.
- The rejecting subarea is *rrr*.
- The transmission group number is *g*.
- The adjacent subarea is *zzz*.
- The explicit route mask is *mmmm*.

Location of failure in a route test

The variable text in message IST523I can help you determine which direction the route test was going when it failed. There are three possibilities. Either the failure is in the adjacent subarea or in the link from the adjacent subarea to the rejecting subarea; or the failure is in the rejecting subarea; or the location of the failure could not be determined. If the reason is "UNEXPECTED TYPE BYTE X'##'", then the location of the failure could not be determined. This condition should not occur.

If the reason is:

- "A REQUIRED TG IS INACTIVE"
- "MIGRATION ER NOT SUPPORTED"
- "MIGRATION NODE DOES NOT SUPPORT THIS ER"
- "MIGRATION NODE ENCOUNTERED"

the adjacent subarea *follows* the rejecting subarea in the route being tested. Therefore, the problem is in the adjacent subarea or the link to the adjacent subarea from the rejecting subarea. (See Figure 32)

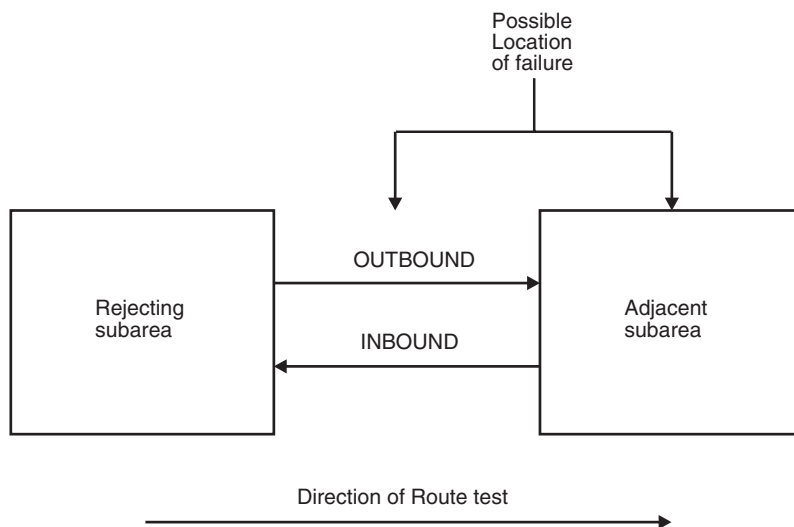


Figure 32. Route test failure (TG inactive or migration)

If the reason is "ER NOT REVERSIBLE," "ER EXCEEDS MAXIMUM LENGTH," or "ER NOT DEFINED," the adjacent subarea *precedes* the rejecting node in the route being tested. (See Figure 33 on page 166.)

Check to see whether the problem is a path definition error. If not, it might be a VTAM error.

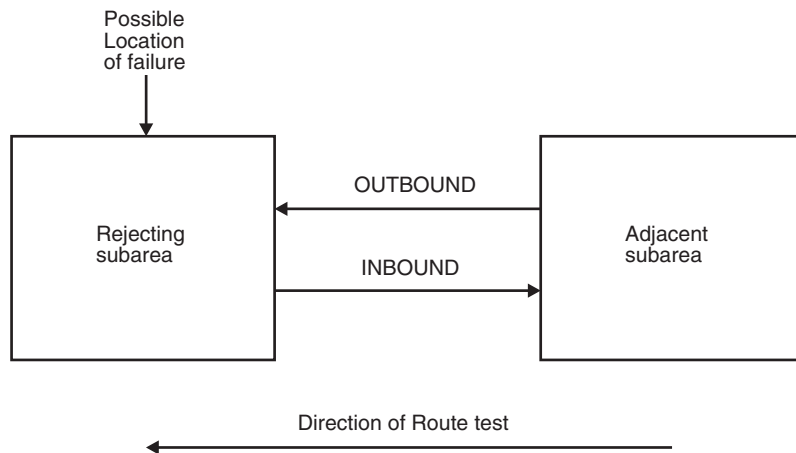


Figure 33. Route test failure (ER not reversible, exceeds maximum length, or not defined)

The ER MASK field indicates which ER numbers the rejecting subarea can use to send data back to the host that requested the test. (These explicit routes are called *reverse explicit routes*.) For example, an ER MASK field of hex 88 (binary 10001000) means that the 0 and 4 bits of the mask are turned on, so ER numbers 0 and 4 can be used to send data back to the host that requested the test. If message IST523I says an explicit route is “NOT REVERSIBLE,” that means the rejecting subarea does not have the reverse explicit routes in its path definitions.

It is possible for the test results for an explicit route to be lost before they are displayed. For example, if a node or link along the explicit route fails between the time VTAM sends the test request and the time VTAM receives the test results, the results will be lost. If this occurs, reenter the DISPLAY ROUTE command for that explicit route.

See z/OS Communications Server: SNA Messages for complete explanations of the messages resulting from the DISPLAY ROUTE command.

Display RTPS options

To find all RTP pipes whose retransmission rate meet or exceed 0.05%, and clear all diagnostic counters after command processing is complete:

```
D NET,RTPS,REXMIT=0.05,CLEAR=ALL
```

See z/OS Communications Server: SNA Operation for the display output.

Tip: The DISPLAY RTPS,REXMIT command has various formats for displaying the RTP pipes information. This sample command does not use any command filters. The DISPLAY RTPS command can be specified with the CPNAME filter to limit the scope of the search.

- Message IST1960I displays information about the RTP pipe with the PU name and the status of the PU.
- Message IST2084I displays the number of pipes displayed that meet or exceed the specified retransmission rate.
- Message IST2248I displays the number of RTP pipes for which all diagnostic counters have been cleared.

Rule: CPNAME is the only command filter allowed with the REXMIT operand.

Using the CLEAR option on D RTPS

To clear all diagnostic counters for all RTPs issue the following command:

```
D NET,RTPS,CLEAR=ALL
```

Message IST2248I displays the number of RTP pipes for which all diagnostic counters have been cleared.

To clear all diagnostic counters for all RTPs destined to *cpname1*:

```
D NET,RTPS,CLEAR=ALL,CPNAME=cpname1
```

Message IST2248I displays all diagnostic counters cleared for all RTP pipes destined to *cpname1*.

To clear all diagnostic counters for a specific RTP pipe:

```
D NET,RTPS,CLEAR=ALL,ID=rtpname
```

Message IST2248I displays all diagnostic counters cleared for the specified RTP pipe.

See Display ID for an RTP PU with HPRDIAG=YES and CLEAR=ALL for more information about cleared diagnostic counters and messages of the cleared diagnostic counters.

Display TDU information

You can use the DISPLAY TOPO command with the LIST=TDUINFO and SCOPE operands to display VTAM topology database update (TDU) processing information that could be used to detect a TDU war in the network. A TDU war is the endless exchange of TDUs in contention over the same topology resource, resulting in continuous performance degradation of the APPN network. The topology resources (nodes and TGs) in contention can be identified and, depending on the nature of the problem, the origin of the TDU war may be isolated. If the TDUDIAG start option is set in all network nodes in the network, you can then use the DISPLAY TOPO command with LIST=TDUDIAG to determine which network nodes are updating the resource sequence numbers (RSNs) of the resources in contention. See DISPLAY TOPO in z/OS Communications Server: SNA Operation for additional information about these commands.

Tip: It is not required that every network node in the network append TDU diagnostic information in TDUs (through the TDUDIAG start option). However, it might not be possible to diagnose the problem if one of the network nodes updating the RSN of the resources in contention during a TDU war has TDUDIAG=NEVER specified or does not have support for TDU diagnostic information in TDUs.

Use the following diagnostic steps when a TDU war is suspected:

1. DISPLAY NET,TOPO,LIST=TDUINFO,SCOPE=ACTIVITY

Use this command to identify the topology resources that are reported in TDUs received and TDUs sent by this node most frequently. In addition, SCOPE=ACTIVITY identifies the topology resources that have had RSNs updated by the host node most frequently, which can indicate whether this node is one of the network nodes involved in a TDU war.

2. DISPLAY NET,TOPO,LIST=TDUINFO,SCOPE=RECENT

This command displays the topology resources reported in TDUs received and TDUs sent by this node most recently. If the resources obtained in the previous step are also in the output generated by this command, observe the following items:

- Resource sequence numbers (RSNs)
- TDU accepted counts (ACC) and TDU rejected counts (REJ) of the resources reported in TDUs received
- TDU sent counts (SENT) and TDU received counts (REC) of the resources reported in TDUs sent

If these displays show repeated receive and send TDUs for the same resource with one of the following symptoms, you might have detected a TDU war:

- Continuous rejection of TDUs received for a resource (TDU rejected count for the resource is rising), with the RSN in the TDUs rising or unchanged
- Continuous acceptance of TDUs received for a resource (TDU accepted count for the resource is rising), with the RSN in the TDUs rising

Note:

- a. There are times that TDUs will flood the network when network nodes propagate topology information to other nodes. For example, when two portions of the same APPN network are connected by CP-CP sessions for the first time, topology information is broadcast in TDUs. The TDU traffic eventually subsides and this is not a TDU war.
- b. When TDUs are continuously sent and the RSN updated for the same resource, but TDUs are never received for that resource with higher RSNs, this might not be a TDU war, but a problem with TGs. You can review the system logs to see whether many error messages have been received for TGs that originate in the node that is sending the TDUs.

3. DISPLAY NET,TOPO,LIST=TDUDIAG

Use this summary command to identify the resources with the most frequent TDU activity (displayed with the DISPLAY NET, TOPO, LIST=TDUINFO,SCOPE=ACTIVITY command) that also contain detailed diagnostic information about RSN updates for that resource.

4. DISPLAY NET,TOPO,LIST=TDUDIAG with the ORIG, DEST, and TGN operands for a TG or with the ID operand for a node

Use these commands to display the detailed diagnostic information about TDUs describing the resources identified with the previous LIST=TDUDIAG summary command. The detailed TDU diagnostic information identifies the network nodes that are updating the RSN of the resource in contention, thus causing the possible TDU war.

Tip: If the TDUDIAG start option is not set in all network nodes in the network that are involved in a TDU war, this information may not be complete and could be misleading. However, the RSN before the RSN update and the RSN after the RSN update are displayed in each TDU diagnostic record, so it is still possible to determine at least one of the nodes involved.

5. DISPLAY NET,TOPO,LIST=TDUINFO,CLEAR or DISPLAY NET,TOPO,LIST=TDUDIAG,CLEAR

Use the CLEAR=YES operand to clear all TDU statistics data and all the TDU diagnostic information collected so far. Subsequent displays can be used to show that TDU activities since the last CLEAR command was issued. In a true TDU war, a large amount of TDU traffic is generated within seconds. Thus, the CLEAR operand is useful to confirm whether a TDU war has started.

Tip: When a TDU war is in progress, the RSN value and TDU counter values for a resource can increase rapidly. An RSN value of `*****` or TDU counter values of `****` in the output from the `DISPLAY NET,TOPO,LIST=TDUINFO` command or the `DISPLAY NET,TOPO,LIST=TDUDIAG` summary command indicates that the values are greater than the available space for those values to be displayed. You can reenter the command with the `FORMAT=LONG` operand to display these values in a format that includes two lines of output for each resource.

Figure 34 and the sample displays that follow show one example of how a TDU war can occur.

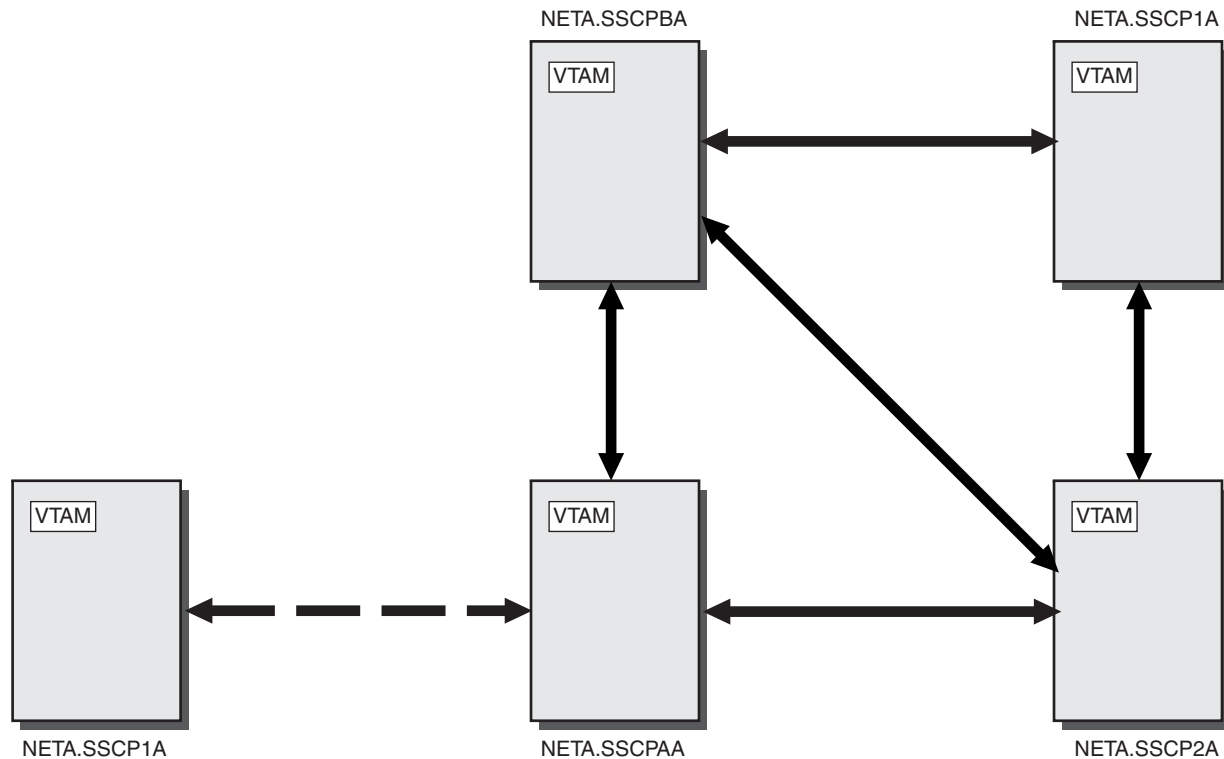


Figure 34. One example of how a TDU war might occur

In Figure 34,

- CP-CP sessions are active between the following nodes:
 - NETA.SSCPBA and NETA.SSCP1A (in the upper right corner)
 - NETA.SSCPBA and NETA.SSCPAA
 - NETA.SSCPBA and NETA.SSCP2A
 - NETA.SSCP2A and NETA.SSCP1A (in the upper right corner)
 - NETA.SSCP2A and NETA.SSCPAA
- No CP-CP sessions are active between NETA.SSCPAA and NETA.SSCP1A (in the lower left corner)

Tip: The default value for the `NUM` operand is used in the following `DISPLAY NET,TOPO,LIST=TDUINFO` and `DISPLAY NET,TOPO,LIST=TDUDIAG` commands to limit the amount of output displayed. You can specify a larger value, up to the maximum value of 50, on the `NUM` operand if an obvious pattern cannot be detected with the default value when you are diagnosing a possible TDU war.

Initially, from NETA.SSCPBA, the DISPLAY NET,TOPO,LIST=TDUINFO command with the SCOPE operand set to ACTIVITY, or allowed to default, produces the following output:

D NET,TOPO,LIST=TDUINFO,SCOPE=ACTIVITY

```

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TDU INFORMATION
IST1780I TOPOLOGY RESOURCES WITH MOST FREQUENT TDU ACTIVITY
IST2275I TDU INFORMATION SINCE LAST RESET ON 02/11/10 AT 10:48:52
IST2290I TDUDIAG START OPTION = 1000
IST2276I NO CORRUPTION OF TOPOLOGY CONTROL VECTORS DETECTED
IST924I -----
IST2286I TDUS RECEIVED:
IST1777I CP NAME          RSN      DESTINATION CP    TGN ACC  REJ
IST1778I NETA.SSCP2A      4        NETA.SSCPBA      21  2    0
IST1778I NETA.SSCPAA     10        NETA.SSCPBA      21  2    0
IST1778I NETA.SSCPBA      4        NETA.SSCP2A      21  0    1
IST1778I NETA.SSCPBA      4        NETA.SSCPAA      21  0    1
IST1778I NETA.SSCP2A      4        NETA.SSCPAA      21  1    0
IST1778I NETA.SSCP2A      4        NETA.SSCP1A      21  1    0
IST1778I NETA.SSCP2A      2        ***NA***         NA  1    0
IST1778I NETA.SSCPAA     12        NETA.SSCP2A      21  1    0
IST1778I NETA.SSCPAA      8        NETA.SSCP1A      21  1    0
IST1778I NETA.SSCPAA      8        ***NA***         NA  1    0
IST2301I 10 OF 13 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2287I TDUS SENT:
IST2288I CP NAME          RSN      DESTINATION CP    TGN SENT REC
IST1778I NETA.SSCPBA      4        NETA.SSCPAA      21  5    1
IST1778I NETA.SSCPBA      4        NETA.SSCP2A      21  4    1
IST1778I NETA.SSCPAA     10        NETA.SSCPBA      21  4    3
IST1778I NETA.SSCP2A      4        NETA.SSCPBA      21  3    3
IST1778I NETA.SSCP2A      4        NETA.SSCPAA      21  3    2
IST1778I NETA.SSCP2A      4        NETA.SSCP1A      21  3    2
IST1778I NETA.SSCP2A      2        ***NA***         NA  3    2
IST1778I NETA.SSCPBA      4        NETA.SSCP1A      21  3    2
IST1778I NETA.SSCPBA      2        ***NA***         NA  3    2
IST1778I NETA.SSCPAA     12        NETA.SSCP2A      21  3    2
IST2301I 10 OF 15 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2289I RESOURCE SEQUENCE NUMBERS UPDATED BY THIS NODE:
IST2292I CP NAME          RSN      DESTINATION CP    TGN  UPDATED
IST2293I NETA.SSCPBA      4        NETA.SSCP2A      21  2
IST2293I NETA.SSCPBA      4        NETA.SSCPAA      21  2
IST2293I NETA.SSCPBA      4        NETA.SSCP1A      21  2
IST2301I 3 OF 3 TOPOLOGY RESOURCES DISPLAYED
IST314I END

```

The DISPLAY NET,TOPO,LIST=TDUINFO command with the SCOPE=RECENT operand specified produces the following output:

D NET,TOPO,LIST=TDUINFO,SCOPE=RECENT

```

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TDU INFORMATION
IST1776I TOPOLOGY RESOURCES WITH MOST RECENT TDU ACTIVITY
IST2275I TDU INFORMATION SINCE LAST RESET ON 02/11/10 AT 10:48:52
IST2290I TDUDIAG START OPTION = 1000
IST2276I NO CORRUPTION OF TOPOLOGY CONTROL VECTORS DETECTED
IST924I -----
IST1779I TDUS RECEIVED BETWEEN 02/11/10 10:49:17 - 02/11/10 10:49:23
IST1777I CP NAME          RSN      DESTINATION CP    TGN ACC  REJ
IST1778I NETA.SSCPBA      2        NETA.SSCP2A      21  0    1
IST1778I NETA.SSCP2A      4        NETA.SSCPBA      21  2    0
IST1778I NETA.SSCP2A      2        NETA.SSCPBA      21  1    0
IST1778I NETA.SSCPBA      2        NETA.SSCPAA      21  0    1

```

```

IST1778I NETA.SSCPAA      10      NETA.SSCPBA      21  2  0
IST1778I NETA.SSCPAA      8       NETA.SSCPBA      21  1  0
IST1778I NETA.SSCP2A      4       NETA.SSCPAA      21  1  0
IST1778I NETA.SSCP2A      4       NETA.SSCP1A      21  1  0
IST1778I NETA.SSCP2A      2       ***NA***        NA  1  0
IST1778I NETA.SSCPAA     12      NETA.SSCP2A      21  1  0
IST2301I 10 OF 15 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2285I TDUS SENT BETWEEN 02/11/10 10:49:17 - 02/11/10 10:49:23
IST2288I CP NAME          RSN      DESTINATION CP    TGN SENT REC
IST1778I NETA.SSCPBA      4       NETA.SSCP2A      21  4  1
IST1778I NETA.SSCP2A      4       NETA.SSCPBA      21  3  3
IST1778I NETA.SSCP2A      2       NETA.SSCPBA      21  2  2
IST1778I NETA.SSCP2A      4       NETA.SSCPAA      21  3  2
IST1778I NETA.SSCP2A      4       NETA.SSCP1A      21  3  2
IST1778I NETA.SSCP2A      2       ***NA***        NA  3  2
IST1778I NETA.SSCPBA      4       NETA.SSCP2A      21  3  0
IST1778I NETA.SSCPBA      4       NETA.SSCPAA      21  5  1
IST1778I NETA.SSCPBA      4       NETA.SSCP1A      21  3  2
IST1778I NETA.SSCPBA      2       ***NA***        NA  3  2
IST2301I 10 OF 49 TOPOLOGY RESOURCES DISPLAYED
IST314I END

```

At this point, CP-CP sessions are activated between SSCPAA and a second network node with the CP name of NETA.SSCP1A (pictured in the lower left corner of Figure 34 on page 169). Because there is an existing network node with this same CP name in the network, this configuration error results in a TDU war. Again from SSCPBA, shortly after CP-CP sessions are activated between SSCPAA and the second NETA.SSCP1A, the DISPLAY NET,TOPO,LIST=TDUINFO,SCOPE=ACTIVITY command produces the following output:

D NET, TOPO, LIST=TDUINFO, SCOPE=ACTIVITY

```

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TDU INFORMATION
IST1780I TOPOLOGY RESOURCES WITH MOST FREQUENT TDU ACTIVITY
IST2275I TDU INFORMATION SINCE LAST RESET ON 02/11/10 AT 10:48:52
IST2290I TDUDIAG START OPTION = 1000
IST2276I NO CORRUPTION OF TOPOLOGY CONTROL VECTORS DETECTED
IST924I -----
IST2286I TDUS RECEIVED:
IST1777I CP NAME          RSN      DESTINATION CP    TGN ACC REJ
IST1778I NETA.SSCP1A      3126    NETA.SSCPAA      21  793 4288
IST1778I NETA.SSCP1A      3128    NETA.SSCPBA      21  793 4260
IST1778I NETA.SSCP1A      3128    NETA.SSCP2A      21  793 4257
IST1778I NETA.SSCPAA      16      NETA.SSCP1A      21  5  2
IST1778I NETA.SSCP2A      4       NETA.SSCPBA      21  2  0
IST1778I NETA.SSCPAA      10      NETA.SSCPBA      21  2  0
IST1778I NETA.SSCPBA      4       NETA.SSCP2A      21  0  1
IST1778I NETA.SSCPBA      4       NETA.SSCPAA      21  0  1
IST1778I NETA.SSCP2A      4       NETA.SSCPAA      21  1  0
IST1778I NETA.SSCP2A      4       NETA.SSCP1A      21  1  0
IST2301I 10 OF 14 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2287I TDUS SENT:
IST2288I CP NAME          RSN      DESTINATION CP    TGN SENT REC
IST1778I NETA.SSCP1A      3126    NETA.SSCPAA      21  5257 5119
IST1778I NETA.SSCP1A      3128    NETA.SSCPBA      21  5231 5101
IST1778I NETA.SSCP1A      3128    NETA.SSCP2A      21  5228 5098
IST1778I NETA.SSCPAA      16      NETA.SSCP1A      21  9  11
IST1778I NETA.SSCPBA      4       NETA.SSCPAA      21  5  1
IST1778I NETA.SSCPBA      4       NETA.SSCP2A      21  4  1
IST1778I NETA.SSCPAA      10      NETA.SSCPBA      21  4  3
IST1778I NETA.SSCP2A      4       NETA.SSCPBA      21  3  3
IST1778I NETA.SSCP2A      4       NETA.SSCPAA      21  3  2

```



```

IST1778I NETA.SSCP2A      4      NETA.SSCP1A      21  3      2
IST2301I 10 OF 16 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2289I RESOURCE SEQUENCE NUMBERS UPDATED BY THIS NODE:
IST2292I CP NAME          RSN      DESTINATION CP    TGN  UPDATED
IST2293I NETA.SSCPBA      4        NETA.SSCP2A      21   2
IST2293I NETA.SSCPBA      4        NETA.SSCPAA      21   2
IST2293I NETA.SSCPBA      4        NETA.SSCP1A      21   2
IST2301I 3 OF 3 TOPOLOGY RESOURCES DISPLAYED
IST314I END

```

The DISPLAY NET,TOPO,LIST=TDUINFO,SCOPE=RECENT command produces the following output:

D NET, TOPO, LIST=TDUINFO, SCOPE=RECENT

```

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TDU INFORMATION
IST1776I TOPOLOGY RESOURCES WITH MOST RECENT TDU ACTIVITY
IST2275I TDU INFORMATION SINCE LAST RESET ON 02/11/10 AT 10:48:52
IST2290I TDUDIAG START OPTION = 1000
IST2276I NO CORRUPTION OF TOPOLOGY CONTROL VECTORS DETECTED
IST924I -----
IST1779I TDUS RECEIVED BETWEEN 02/11/10 13:40:14 - 02/11/10 13:40:14
IST1777I CP NAME          RSN      DESTINATION CP    TGN  ACC  REJ
IST1778I NETA.SSCP1A      3126    NETA.SSCPAA      21   793  4288
IST1778I NETA.SSCP1A      3128    NETA.SSCP2A      21   793  4257
IST1778I NETA.SSCP1A      3124    NETA.SSCPAA      21   792  4288
IST1778I NETA.SSCP1A      3122    NETA.SSCPAA      21   791  4288
IST1778I NETA.SSCP1A      3128    NETA.SSCPBA      21   793  4260
IST1778I NETA.SSCP1A      3126    NETA.SSCP2A      21   792  4257
IST1778I NETA.SSCP1A      3126    NETA.SSCPBA      21   792  4260
IST1778I NETA.SSCP1A      3124    NETA.SSCPBA      21   791  4260
IST1778I NETA.SSCP1A      3124    NETA.SSCP2A      21   791  4257
IST1778I NETA.SSCP1A      3122    NETA.SSCP2A      21   790  4257
IST2301I 10 OF 50 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2285I TDUS SENT BETWEEN 02/11/10 13:40:14 - 02/11/10 13:40:20
IST2288I CP NAME          RSN      DESTINATION CP    TGN  SENT  REC
IST1778I NETA.SSCP1A      3126    NETA.SSCPAA      21   5257  5119
IST1778I NETA.SSCP1A      3128    NETA.SSCP2A      21   5228  5098
IST1778I NETA.SSCP1A      3124    NETA.SSCPAA      21   5256  5118
IST1778I NETA.SSCP1A      3122    NETA.SSCPAA      21   5255  5117
IST1778I NETA.SSCP1A      3128    NETA.SSCPBA      21   5231  5101
IST1778I NETA.SSCP1A      3126    NETA.SSCP2A      21   5227  5097
IST1778I NETA.SSCP1A      3126    NETA.SSCPBA      21   5230  5100
IST1778I NETA.SSCP1A      3124    NETA.SSCPBA      21   5229  5099
IST1778I NETA.SSCP1A      3124    NETA.SSCP2A      21   5226  5096
IST1778I NETA.SSCP1A      3122    NETA.SSCP2A      21   5225  5095
IST2301I 10 OF 50 TOPOLOGY RESOURCES DISPLAYED
IST314I END

```

These last two commands display the following symptoms of a TDU war:

- There are many inbound and outbound TDUs describing the same resources.
- The resource sequence number (RSN) and TDU rejection count (REJ) in message IST1778I for the following resources increase continuously:
 - TG oriented from NETA.SSCP1A to NETA.SSCPBA with TG number 21
 - TG oriented from NETA.SSCP1A to NETA.SSCP2A with TG number 21
 - TG oriented from NETA.SSCP1A to NETA.SSCPAA with TG number 21
- Because NETA.SSCPBA continuously receives TDUs about the above three resources, they become the most active resources (in terms of TDUs received) reported by DISPLAY NET,TOPO,LIST=TDUINFO,SCOPE=ACTIVITY command.

If the TDUDIAG start option is set in one or more of the network nodes involved in the TDU war, additional TDU diagnostic information will be collected about the network nodes that are updating the RSNs of these resources. The DISPLAY NET,TOPO,LIST=TDUDIAG command produces the following output:

D NET, TOPO, LIST=TDUDIAG

```

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TDU DIAGNOSTICS
IST2274I TDU DIAGNOSTIC SUMMARY:
IST1780I TOPOLOGY RESOURCES WITH MOST FREQUENT TDU ACTIVITY
IST2308I THAT HAVE SAVED TDUDIAG RSN UPDATES
IST2275I TDU INFORMATION SINCE LAST RESET ON 02/11/10 AT 13:11:32
IST2290I TDUDIAG START OPTION = 1000
IST2276I NO CORRUPTION OF TOPOLOGY CONTROL VECTORS DETECTED
IST924I -----
IST2286I TDUS RECEIVED:
IST1777I CP NAME          RSN      DESTINATION CP    TGN ACC  REJ
IST1778I NETA.SSCP1A      3126    NETA.SSCPAA      21  793  4288
IST1778I NETA.SSCP1A      3128    NETA.SSCP2A      21  793  4257
IST1778I NETA.SSCP1A      3128    NETA.SSCPBA      21  793  4260
IST2301I 3 OF 3 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2287I TDUS SENT:
IST2288I CP NAME          RSN      DESTINATION CP    TGN SENT REC
IST1778I NETA.SSCP1A      3126    NETA.SSCPAA      21  5257  5119
IST1778I NETA.SSCP1A      3128    NETA.SSCP2A      21  5228  5098
IST1778I NETA.SSCP1A      3128    NETA.SSCPBA      21  5231  5101
IST2301I 3 OF 3 TOPOLOGY RESOURCES DISPLAYED
IST924I -----
IST2289I RESOURCE SEQUENCE NUMBERS UPDATED BY THIS NODE:
IST2301I 0 OF 0 TOPOLOGY RESOURCES DISPLAYED
IST314I END

```

Additional detailed TDU diagnostic information about the network nodes that are updating the RSNs of each TG can be displayed with LIST=TDUDIAG. The DISPLAY NET,TOPO,LIST=TDUDIAG,ORIG=SSCP1A,DEST=SSCPBA,TGN=21 command produces the following output:

D NET, TOPO, LIST=TDUDIAG, ORIG=SSCP1A, DEST=SSCPBA, TGN=21

```

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TDU DIAGNOSTICS
IST2311I TDU DIAGNOSTIC INFORMATION FOR TG: TGN = 21
IST2256I ORIG = NETA.SSCP1A - DEST = NETA.SSCPBA
IST2312I CURRENT RSN = 3128 - HEX RSN = 00000C38
IST924I -----
IST2275I TDU INFORMATION SINCE LAST RESET ON 02/11/10 AT 10:48:52
IST1769I LAST TDU RECEIVED - 02/11/10 13:40:14 FROM NETA.SSCP1A
IST2281I LAST TDU SENT - 02/11/10 13:40:20
IST2282I TDU COUNTS:
IST2352I SENT = 3890 RECEIVED = 5101
IST2353I ACCEPTED = 793 REJECTED = 4260
IST2354I IGNORED = 48
IST2313I RSN UPDATE COUNT = 4260
IST924I -----
IST2294I TDUDIAG RSN UPDATES:
IST2295I TIME HEX RSN HEX RSN
IST2296I CP NAME UPDATED BEFORE AFTER REASON
IST2297I NETA.SSCP1A 13:40:20 00000C36 00000C38 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP1A
IST2297I NETA.SSCP1A 13:40:20 00000C34 00000C36 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP2A
IST2297I NETA.SSCP1A 13:40:19 00000C32 00000C34 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP1A
IST2297I NETA.SSCP1A 13:40:19 00000C30 00000C32 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCPAA

```

```

IST2297I NETA.SSCP1A      13:40:19 00000C2E 00000C30 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP1A
IST2297I NETA.SSCP1A      13:40:19 00000C2C 00000C2E TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCPAA
IST2297I NETA.SSCP1A      13:40:19 00000C2A 00000C2C TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP1A
IST2297I NETA.SSCP1A      13:40:19 00000C28 00000C2A TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP2A
IST2297I NETA.SSCP1A      13:40:19 00000C26 00000C28 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP1A
IST2297I NETA.SSCP1A      13:40:18 00000C24 00000C26 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCPAA
IST2314I 10 OF 50 RSN UPDATES DISPLAYED
IST314I END

```

Because the previous command was entered on SSCPBA, which is not the origin (owner) of the TG, all of the RSN updates are inbound to SSCPBA. Therefore, you cannot tell if the RSN for the resource is being updated by one network node with a CP name of NETA.SSCP1A, or two. When the RSN increased by two and all of the TDUDIAG RSN updates display the same CP name in message IST2297I, it is usually an indication of two network nodes with the same CP name. To determine whether this is the case, you can enter the same command on the NETA.SSCP1A known to you, which produces the following output:

D NET, TOPO, LIST=TDUDIAG, ORIG=SSCP1A, DEST=SSCPBA, TGN=21

```

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TDU DIAGNOSTICS
IST2311I TDU DIAGNOSTIC INFORMATION FOR TG: TGN = 21
IST2256I ORIG = NETA.SSCP1A - DEST = NETA.SSCPBA
IST2312I CURRENT RSN = 3128 - HEX RSN = 00000C38
IST2355I TDUDIAG THRESHOLD REACHED ON 02/11/10 AT 13.39.00
IST924I -----
IST2275I TDU INFORMATION SINCE LAST RESET ON 02/11/10 AT 10:48:52
IST1769I LAST TDU RECEIVED - 02/11/10 13:40:14 FROM NETA.SSCP1A
IST2281I LAST TDU SENT - 02/11/10 13:40:20
IST2282I TDU COUNTS:
IST2352I SENT = 793 RECEIVED = 793
IST2353I ACCEPTED = 0 REJECTED = 793
IST2354I IGNORED = 0
IST2313I RSN UPDATE COUNT = 793
IST924I -----
IST2294I TDUDIAG RSN UPDATES:
IST2295I
IST2296I CP NAME          TIME      HEX RSN  HEX RSN  REASON
IST2297I NETA.SSCP1A      13:40:20 00000C36 00000C38 TDU GREATER
IST2297I NETA.SSCP1A      13:40:20 00000C34 00000C36 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP2A
IST2297I NETA.SSCP1A      13:40:19 00000C32 00000C34 TDU GREATER
IST2297I NETA.SSCP1A      13:40:19 00000C30 00000C32 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCPBA
IST2297I NETA.SSCP1A      13:40:19 00000C2E 00000C30 TDU GREATER
IST2297I NETA.SSCP1A      13:40:19 00000C2C 00000C2E TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCPBA
IST2297I NETA.SSCP1A      13:40:19 00000C2A 00000C2C TDU GREATER
IST2297I NETA.SSCP1A      13:40:19 00000C28 00000C2A TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCP2A
IST2297I NETA.SSCP1A      13:40:19 00000C26 00000C28 TDU GREATER
IST2297I NETA.SSCP1A      13:40:18 00000C24 00000C26 TDU GREATER
IST2300I RECEIVED FROM: NETA.SSCPBA
IST2314I 10 OF 50 RSN UPDATES DISPLAYED
IST314I END

```

Messages IST2297I and IST2300I are displayed for inbound TDUs. For outbound TDUs, only message IST2297I is displayed. From this display, you can see that

TDUs with RSN updates are both inbound to and outbound from this node. If the TDU is inbound, the RSN was not updated by this node. If the TDU is outbound, the RSN was updated by this node. This pattern indicates that there is another network node with a CP name of NETA.SSCP1A in the network.

Display traces

Use the DISPLAY TRACES command to display the status of a trace for a resource and its subordinate nodes. DISPLAY TRACES applies to the following types of traces:

BUF	Buffer contents trace
CNM	Communication network management trace
EXIT	Session management exit (SME) buffer trace
GPT	Generalized PIU trace
IO	Input/output trace
LINE	NCP line trace
MODULE	Module trace
NETCTLR	3710 network controller trace
QDIOSYNC	Queued Direct I/O Diagnostic Synchronization
ROUTE	APPN route selection trace
SIT	Scanner interface trace
SMS	Storage management services buffer use trace
STATE	Resource state trace
TG	Transmission group trace
TSO	TSO user ID trace
VTAM	VTAM internal trace (VIT)

Display VTAM storage

Use the DISPLAY VTAMSTOR command to display storage contents associated with:

- Storage address
- VTAM module
- Network address
- VTAM resource

Display workload information for a device

Storage problems can be related to a specific I/O device. Because outbound data cannot be transmitted to an I/O device until it is accepted by that device (that is,

until write processing completes), there are scenarios in which this storage associated with this data can accumulate at the DLC (data link control) layer. For all I/O devices that perform real I/O that are represented by a TRLE (predefined or dynamically built), VTAM tracks the outbound workload (units of work) for each device. This tracking mechanism allows console operators to isolate this type of problem to a specific device.

The console operator can quickly isolate a storage problem to a specific device using the DISPLAY TRL command. If a device has exceeded internal thresholds, message IST1800I is issued with the text **** CONGESTED ****. Additional details regarding the workload for a specific device are displayed with the DISPLAY TRL,TRLE=*trlename* command. VTAM displays the current, average, and the maximum workload for each device. When the current workload is excessive, the I/O activity for this device might be associated with system storage shortages.

If the counts for a device reveal an excessive current workload, additional steps are required to isolate the problem.

- **Steps for a console operator**

When a device is marked as congested, further action is required to determine whether the congestion is related to a system storage problem. If the following steps indicate that a system storage shortage is present, it might be necessary to obtain documentation (such as a console log and a dump) to diagnose the congestion related to this device. This condition might be relieved by deactivating the PU (or stopping the device for TCP/IP).

1. Review the system console for any messages related to current storage shortage conditions.
2. Issue the following VTAM display commands:
 - D NET,CSM
 - D NET,BFRUSE
 - D NET,STORUSE,POOL=*

Note: If applicable, also issue the TCP/IP DISPLAY command D TCPIP,,STOR.

3. Issue D NET,TRL,TRLE=*trlename* to obtain more details about the device congestion. Message IST1802I displays detailed counts of units of work for the device measured at the Data Link Control (DLC) layer.
4. Activate VTAM tuning statistics (TNSTAT), RMF™, or other monitoring tools to monitor this specific device.
5. Display the active jobs in the system to determine whether new work was recently started.

- **Steps for a system programmer**

The following steps might be required to isolate a system storage problem that is related to an I/O device:

1. Review the network configuration related to this device or any recent configuration changes for this system.
2. Review or monitor (using the output from VTAM TNSTAT or RMF) the network traffic related to this device. Compare the actual workload to the I/O capacity of the hardware device.
3. Determine if the congestion is related to a specific time of day, job, application, or type of workload.
4. Verify that missing interrupt handler (MIH) is enabled for the write devices.

5. Review or verify that the maintenance level for the hardware device is current.
6. Consider automating the necessary storage displays to monitor system conditions.

Using VTAM MODIFY commands for problem determination

This topic includes the following tasks:

- “Issuing the MODIFY CSDUMP command”
- “Modifying input/output problem determination”
- “Modifying message module identification” on page 178
- “Modifying NCP intensive mode recording” on page 179
- “Modifying SDLC link level 2 test” on page 179
- “Issuing the MODIFY TOPO command to clear EE connection network unreachable partner information” on page 180
- “Modifying tuning statistics” on page 181
- “Issuing the MODIFY VTAMOPTS command to change start option values” on page 181

See z/OS Communications Server: SNA Operation for additional information about the VTAM commands.

Issuing the MODIFY CSDUMP command

Issue the MODIFY CSDUMP command to do the following tasks:

- Immediately dump the current address space
- Set up a trigger that starts a dump of the current address space when a particular sense code is issued
- Set up a trigger that starts a dump of the current address space when a particular message is issued
- Delete active message or sense code triggers

Tip: You can also use the CSDUMP start option to set the CSDUMP message and sense code trigger.

Modifying input/output problem determination

Use the input/output problem determination (IOPD) facility to detect pending I/O requests when VTAM sends a request to another part of the network and no response is received after a certain period of time. The IOINT start option determines the length of time during which a response must be received.

You can perform the following tasks with the IOPD facility:

- Enable the IOPD facility
- Change the value of the IOINT start option
- Instruct the IOPD facility to write just one message group for each type of pending I/O operation, rather than one group for each operation

To enable the IOPD facility, issue the MODIFY IOPD command or the MODIFY VTAMOPTS command, or set the IOINT start option.

To change the value of the IOINT start option, issue the MODIFY VTAMOPTS command.

- For more information about the MODIFY VTAMOPTS command, see “Issuing the MODIFY VTAMOPTS command to change start option values” on page 181.
- For more information about the MODIFY IOPD command, see z/OS Communications Server: SNA Operation.

During the initialization of a large VTAM network, you might see more pending I/O operations than usual. If you use the IOPD facility to track I/O problems during initialization, the number of message groups issued can degrade your network's performance.

To instruct the IOPD facility to write just one message group for each type of pending I/O operation, rather than one group for each operation, use the IOMSGLIM start option. The resulting reduction in the number of messages issued can improve your network's performance during initialization.

The IOPD facility issues messages IST530I or IST1278I, IST1051I, and IST1062I for each operation that is pending longer than the specified time interval. See z/OS Communications Server: SNA Messages for a description of these messages. For more information about event codes and event IDs, see z/OS Communications Server: IP and SNA Codes.

These messages are only an indication that a problem might exist. The longer an operation remains pending (for example, the more messages issued for the same request unit), the more likely it is that a problem exists. See “Wait” on page 61 for more information about identifying pending I/O problems.

Modifying message module identification

You can choose to include in VTAM messages the last 5 characters of the VTAM module that issued the message. The module name abbreviation is displayed between the message ID and the message text.

To insert or delete module name abbreviations, use one of the following methods, where YES indicates to insert the abbreviations and NO indicates to not include them:

- Specify MSGMOD=YES|NO in the start option
- Issue the F net,MSGMOD=YES|NO command
- Issue the F net,VTAMOPTS,MSGMOD=YES|NO command

MSGMOD=NO is the default value. For more information about the MODIFY MSGMOD and MODIFY VTAMOPTS commands, see z/OS Communications Server: SNA Operation.

Examples:

- If you specify MSGMOD=YES, VTAM message xxxxx, where xxxxx is an operating system unique message number, is displayed as:
xxxxx INFXI DUMP OF ncpname COMPLETE
- If you specify MSGMOD=NO, the message is displayed as:
xxxxx DUMP OF ncpname COMPLETE

Note:

1. Any message that exceeds the maximum message length after the insertion of the module ID is truncated.
2. If your installation has changed the message text and omitted the message ID, the module name is the first item in the message.

Modifying NCP intensive mode recording

You can use intensive mode error recording to record and signal each temporary error over a link and its cause to VTAM.

When you receive many temporary errors for a line, but not enough to create a permanent error, you can use intensive mode recording to find the cause of the error.

To start NCP intensive mode recording, issue the MODIFY IMR command.

The MODIFY IMR command causes the owning system services control point (SSCP) to send a request to the NCP. The NCP then builds and sends RECMS RUs to the SSCP each time an error occurs. The SSCP writes these error records on the LOGREC file. It can also optionally pass them to a user-defined communication network management (CNM) application program, such as the NetView program. See MODIFY IMR command in *z/OS Communications Server: SNA Operation* for more information.

Modifying SDLC link level 2 test

To test NCP Synchronous Data Link Control (SDLC) links, use the SDLC link level 2 (LL2) test.

If you run this test over an extended period of time, you can increase the possibility of repeating an intermittent error that is hard to re-create.

Rule: If you want to test connectivity to a physical unit only, activate the physical unit instead of using an LL2 test.

Depending on the SDLC link you want to test, do one of the following tasks:

- To test an SDLC link between an NCP and a physical unit that is attached on a multipoint line:
 1. Ensure that the physical unit name in the ID operand is inactive and dedicated to the test. Other physical units on the same link can remain active.
 2. Issue the MODIFY LL2 command.
- To test an SDLC link between two NCPs:
 1. Ensure that the primary link station is inactive. The primary link station is the link named in the LL2 command.
 2. Ensure that the secondary link station is active. The secondary link station is the link that will respond to the test.
 3. Issue the MODIFY LL2 command. Start the test from the primary link station.

The MODIFY LL2 command causes the SSCP to send a test RU to the NCP to which the test terminal is connected. The NCP returns test results to the requesting SSCP in a Record Test Results RU.

The test results include the following information:

- The number of test frames transmitted by the node specified by the *nodename* parameter
- The number of test frames received by the node specified by the *nodename* parameter
- The number of test frames received without error by the node specified by the *nodename* parameter (test frames that were successfully returned by the test station)

If these three numbers are not the same, an error in the link or a physical unit has occurred. Use a line trace to further isolate the problem.

To test the SDLC link, VTAM sends test data over the link from the controlling NCP to the remote station, which is an NCP or a peripheral physical unit. The data is then echoed back to the sending NCP. This NCP then compares the data received with the data sent and forwards the results to VTAM. When VTAM receives the test results, it sends message IST549I to the initiating console to indicate how the link level 2 test ended: data returned without errors, data returned with errors, or no data returned because of an inoperative link or initialization error. See message IST549I in *z/OS Communications Server: SNA Messages* for a complete message description.

Issuing the MODIFY TOPO command to clear EE connection network unreachable partner information

To manually clear Enterprise Extender (EE) connection network unreachable partner information, issue the MODIFY TOPO command. This action can make the unreachable paths available for route selection after underlying connection problems are corrected.

To indicate from which network nodes the unreachable partner paths are to be cleared, use the SCOPE operand.

- The default value, *SCOPE=LOCAL*, indicates that the unreachable partner paths are to be cleared from only the network node on which the command is entered.
- A value of *SCOPE=NETWORK* indicates that the unreachable partner paths are to be cleared from all network nodes in the network.

To clear EE connection network unreachable partner information under different conditions, issue the following commands:

- To clear EE connection network unreachable partner information that is associated with a specific virtual node from all network nodes in the network, issue the following command:

```
F procname,TOPO,FUNCTION=CLRUNRCH,VRN=cp_name,SCOPE=NETWORK
```

- To clear EE connection network unreachable partner information that is associated with a specific origin node from only the network node on which the command is entered, issue the following command:

```
F procname,TOPO,FUNCTION=CLRUNRCH,ORIG=cp_name
```

- To clear EE connection network unreachable partner information that is associated with a specific partner node from only the network node on which the command is entered, issue the following command:

```
F procname,TOPO,FUNCTION=CLRUNRCH,DEST=cp_name
```


To control the scope of the unreachable partner information that is cleared, use the ORIG, VRN, and DEST operands in any combination.

Modifying tuning statistics

You can perform the following tasks to modify tuning statistics:

- To initiate recording of tuning statistics, use the TNSTAT start option or issue the MODIFY TNSTAT operator command.
- To stop or start the recording of tuning statistics, or to adjust the tuning statistics controls at any time, issue the MODIFY NOTNSTAT or MODIFY TNSTAT command.
- To initiate recording for all devices (global TNSTATs), specify the VTAM TNSTAT start option.
- To initiate global TNSTATs, issue the MODIFY TNSTAT command with the ACTION=ACTIVATE operand (or allow it to have default settings) and without the TRLE operand.
- To stop global TNSTATs, issue the MODIFY NOTNSTAT command without the TRLE operand.
- To determine which devices are recording or what the tuning statistics controls are set to, issue the DISPLAY TNSTAT command.
- If you do not want to modify recording on all devices, specify the TRLE operand on both MODIFY commands to initiate and stop recording on a TRLE basis.

By specifying the TRLE operand, recording is modified for only those devices managed by the TRLEs that you specify. If you do not specify the TRLE operand, all devices are modified.

The CNSL and TIME values apply to all devices that are actively recording. These values are unaffected by the presence or lack of the TRLE operand.

Summary records can be written to SMF and the system console. If SMF is available, records are written to SMF. If the CNSL TNSTAT parameter is set to YES, summary records are sent to the system console.

Issuing the MODIFY VTAMOPTS command to change start option values

Issue the MODIFY VTAMOPTS command to change certain values that are specified on VTAM start options.

For a description of start options that you can change by issuing this command, see *z/OS Communications Server: SNA Operation*.

Chapter 5. Using dumps

This topic covers the dumps that you can use for problem determination for the VTAM program. The included dumps are:

- MVS Dumps
 - “Abend dump”
 - “Coupling facility structures dump” on page 184
 - “FFST dump” on page 184
 - “Stand-alone dump” on page 184
 - “SVC dump” on page 184
- “Network control program (NCP) dump” on page 185
- “Communication scanner processor dump (3720, 3725, and 3745 only)” on page 187
- “Maintenance and operator subsystem dump (3720, 3725, and 3745 only)” on page 188

“Formatting and printing dump output” on page 188 describes the service aids available for formatting and printing dump output.

For information on dumps generated by First Failure Support Technology™ (FFST), see z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT.

Dumps on MVS operating system

Several kinds of dumps can be produced in an MVS system, depending on the type of failure and operator action:

- “Abend dump”
- “Coupling facility structures dump” on page 184
- “FFST dump” on page 184
- “Stand-alone dump” on page 184
- “SVC dump” on page 184

Abend dump

If the appropriate DD card exists, an abend dump is produced when one of the following conditions occurs:

- The operator enters a CANCEL command.
- An abend macroinstruction is issued.
- A job abnormally ends.

To get an abend dump, the input stream for VTAM must contain a DD statement with the ddname SYSUDUMP or SYSABEND. The resulting dump is written to the data set specified on the SYSUDUMP or SYSABEND DD card. The contents of the dump depend on user specifications. See Table 48 on page 649 to determine what document has more information on the abend dump.

Coupling facility structures dump

When using GR, MNPS, TSO/GR, TCP/IP Sysplexports, or TCP/IP Sysplex Wide Security Associations, you should also dump the coupling facility structures involved when documenting problems with those functions. See Table 48 on page 649 to determine what document has more information on the coupling facility structures dump.

FFST dump

For information on dumps generated by First Failure Support Technology (FFST), see *z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT*.

Stand-alone dump

The stand-alone dump is produced when the operator invokes the stand-alone dump program. This program can be invoked when the operating system is in a disabled wait state or looping. The stand-alone dump may be a high-speed dump, which is not formatted, or a low-speed dump, which is formatted. The dump output is written to the tape or printer (low-speed only) specified on the output operands. The output for a high-speed dump can be formatted for viewing by IPCS. See *z/OS Information Roadmap* to determine what document contains more information on the stand-alone dump.

SVC dump

SVC dumps are produced under these conditions:

- VTAM produces an SVC dump automatically when a program exception occurs. VTAM might be terminated as part of this process. The system log indicates the location of the dump output and whether the dump was successful.
- An operator issues the MODIFY CSDUMP command without specifying any triggers, which causes an immediate dump. This action does not cause VTAM to stop.
- An operator uses the CSDUMP start option or the MODIFY CSDUMP command to set the dump triggers, and an event occurs that matches one of the triggers to take the dump. This action does not cause VTAM to stop.
- An operator can request a dump with the operating system DUMP command. This will not cause VTAM to stop.
- An operator uses a SLIP command with ACTION=SVCD specified, and an event occurs that matches the trap indicated in the SLIP.
- A macroinstruction issues an abend, and there is a DD statement with ddname=SYSMDUMP.
- An SDUMP macroinstruction is issued.
- System recovery routines produce an SVC dump if VTAM causes an error, such as a program exception or abend.

An SVC dump is written to a SYS1.DUMPnn data set (if allocated), the SYSMDUMP output data set, or the data set specified on the DCB operand of the SDUMP macroinstruction. An SVC dump can be formatted for viewing by IPCS. To determine what document contains more information on SVC dump, see *z/OS Information Roadmap*.

Network control program (NCP) dump

You can get either dynamic or static dumps of NCP storage in a communication controller.

You can use the DISPLAY NCPSTOR command to dynamically dump up to 256 bytes of NCP storage. When you request a dynamic dump, VTAM sends repeated DISPSTOR RUs to the NCP until the entire contents of NCP storage have been transmitted. The NCP continues to operate during this period. As a result, the dump represents NCP storage over a period of time.

You can get a static dump with the VTAM dump facility (for a channel-attached or link-attached communication controller). You can also get a static dump with the independent dump utility (for a channel-attached communication controller only). Table 12 on page 186 summarizes the methods and requirements for dumping the NCP.

NCP dumps are not allowed on lines or data links (including multipoint subarea links) for which IPL=NO was specified in the NCP definition deck. Therefore:

- Code IPL=YES for any line or link which might be used to IPL or dump any attached type 4 physical unit.
- If IPL=YES was not specified for a line with one or more type 4 physical units attached, do not try to load or dump over that line.

Note: If you do not specify a value for IPL, the default is IPL=NO.

If you have a 3720 or 3745 Communication Controller with a hard disk in your network, you can transfer NCP, MOSS, and CSP dumps from the communication controller disk to the host. You can also transfer the NCP load module from the host to the disk. If you need to restart the NCP, you can load the NCP load module from the disk.

Specifically, you can perform the following tasks with the 3720 or 3745 Communication Controller with hard disk:

- save the NCP load module to the disk
- load the NCP load module from the disk
- transfer an NCP, MOSS, or CSP dump stored on the disk to the host
- purge an NCP, MOSS, or CSP dump stored on the disk
- control automatic loading and dumping of the NCP to or from the disk
- display up to 256 bytes of NCP dump or state vector stored on the disk

When to use the NCP dump

A dump of the NCP should be taken whenever the NCP abnormally terminates or when an error is suspected in the NCP. It may be possible to determine that a problem exists in the NCP by using the VTAM I/O trace to determine what PIUs are being sent to and received from the communication controller and by using the NCP line trace to determine what is happening on the lines between the communication controller and the link-attached logical unit.

Using the NCP option on MODIFY DUMP

To dump the NCP, use the MODIFY DUMP command with TYPE=NCP specified. For more information on this command, see z/OS Communications Server: SNA Operation. You must execute a separate job to print the dump.

Note: A data host cannot load or dump an NCP.

You can use the DUMPDS operand of the MODIFY DUMP command to specify the file into which you want the dump transferred. If you omit this operand, VTAM uses the file specified on the PCCU definition statement for the NCP.

Note: If you use the same file on the MODIFY DUMP command as was named on the PCCU definition statement for the NCP, an earlier NCP dump may be overwritten.

If you omit the DUMPSTA operand, VTAM uses the link station specified on the VARY ACT command or the PCCU definition statement for the NCP, in that order.

If you are using SNA network interconnection, and you want a gateway NCP to perform a dump, you can set its link station name on the DUMPSTA operand only if that NCP is in the same network as the host processor requesting the dump.

If you set AUTODUMP=YES on the PCCU definition statement, a dump is taken automatically if the NCP abnormally terminates. This dump is written to the dump file named on the DUMPDS option of the PCCU definition statement. If AUTOIPL=YES is specified on the PCCU definition, the NCP is automatically reloaded after the dump is taken. With the 3720 or 3745 Communication Controller with hard disk, you can specify DUMPLOAD=YES on the VARY ACT command, which automatically stores the dump on the disk and loads the NCP load module from the disk.

Note: For the 3705 only, the NCP is partially overwritten in the communication controller storage when the dump is taken. If the NCP is running with partitioned emulation programming (PEP), the emulation routines are included in the dump and are also partially overwritten. The 3705 must be reactivated after the dump is taken.

Table 12. Dumping the NCP

Method of starting NCP dump	Channel or link attached	Requirements	NCP status	Printing dump
MODIFY DUMP command, OPTION=DYNA	Both	PCCU definition requirements: DUMPDS must be specified.	Active throughout dump, no reactivation required.	Execute SSP dump formatting program IFLDUMP
MODIFY DUMP command, OPTION=STATIC	Both	PCCU definition requirements: DUMPDS must be specified.	Deactivated when dump is completed; operator must reactivate the NCP.	Execute SSP dump formatting program IFLDUMP

Table 12. Dumping the NCP (continued)

Method of starting NCP dump	Channel or link attached	Requirements	NCP status	Printing dump
VTAM Error Recovery Procedures (automatic)	Both	PCCU definition requirements: DUMPDS must be specified. AUTODMP=YES must be specified. AUTOIPL may be specified.	Activated after dump is completed if AUTOIPL=YES and restart is successful, or if AUTOIPL=NO and operator requests communication controller IPL.	Execute SSP dump formatting program IFLDUMP
VTAM Error Recovery Procedures (with operator intervention)	Both	PCCU definition requirements: DUMPDS must be specified. AUTODMP=NO must be specified. AUTOIPL may be specified.	Activated after dump is completed if AUTOIPL=YES and restart is successful, or if AUTOIPL=NO and operator requests communication controller IPL.	Execute SSP dump formatting program IFLDUMP
NCP dump utility: IFLREAD	Channel-attached only	See z/OS Information Roadmap to determine what document describes the NCP dump utility.	Deallocate communication controller from VTAM and allocate to the independent dump utility.	IFLREAD uses IFLDUMP to print the dump automatically.
Controller-detected error (3720 or 3745 with disk only)	Both	VARY ACT requirements: DUMpload=YES must be specified. The dump slot for this central control unit (CCU) on the controller's hard disk must be empty.	Activated after dump is completed if AUTOIPL=YES and restart is successful, or if AUTOIPL=NO and operator requests communication controller IPL.	Execute SSP dump formatting program IFLDUMP Transfer to host using MODIFY DUMP with ACTION=TRANSFER.

Using the independent NCP dump utility (channel-attached controller only)

To use the independent NCP dump utility, the communication controller must be inactive. See z/OS Information Roadmap to determine what document describes the following information:

- Job control language needed to invoke the independent NCP dump utility
- NCP data areas, registers, and codes found in an NCP storage dump

Communication scanner processor dump (3720, 3725, and 3745 only)

The communication scanner processor (CSP) automatically dumps its contents when it detects an error. CSP stores the dump on the communication controller disk, or on the MOSS diskette, and the NCP sends an alert message to the host to inform it of the error. You can use the MODIFY DUMP command to transfer this

dump to a dump file in the host processor. After the dump has been transferred to the host, you must then run a separate job to print the dump.

If you have a 3720 or 3745 Communication Controller with hard disk in your network, you can also purge the CSP dump from the communication controller with the MODIFY DUMP command.

To transfer the contents of the CSP dump, use the MODIFY DUMP command with TYPE=CSP specified. You can use the DUMPDS operand of the MODIFY DUMP command to specify the file into which you want the dump transferred. If you omit this operand, the dump is put into one of two dump files:

- The file specified by the CDUMPDS operand of the PCCU definition statement for the NCP
- The file specified by the DUMPDS operand of the PCCU definition statement for the NCP (if CDUMPDS is not specified on the PCCU definition)

Maintenance and operator subsystem dump (3720, 3725, and 3745 only)

The maintenance and operator subsystem (MOSS) automatically dumps its contents when it detects an error. MOSS stores the dump on the communication controller disk, or on the MOSS diskette, and the NCP sends an alert message to the host to inform it of the error. You can use the MODIFY DUMP command to transfer this dump to a dump file in the host processor. After the dump has been transferred to the host, you must then run a separate job to print the dump.

If you have a 3720 or 3745 Communication Controller with hard disk in your network, you can also purge the MOSS dump from the communication controller with the MODIFY DUMP command.

To transfer the contents of the MOSS dump, use the MODIFY DUMP command with TYPE=MOSS specified. You can use the DUMPDS operand of the MODIFY DUMP command to specify the file into which you want the dump transferred. If you omit this operand, the dump is put into one of two dump files:

- The file specified by the MDUMPDS operand of the PCCU definition statement for the NCP
- The file specified by the DUMPDS operand of the PCCU definition statement for the NCP (if MDUMPDS is not specified on the PCCU definition).

Formatting and printing dump output

The service aids described in this information are available for formatting and printing dump output.

IPCS service aids

IPCS processes SVC dumps and high-speed stand-alone dumps for online viewing. For information on using IPCS with VTAM, see Chapter 6, "Using VTAM dump analysis tools," on page 191.

To determine what document further describes IPCS commands, see z/OS Information Roadmap.

ABDUMP service aid

ABDUMP operates as part of the operating system abnormal termination (abend) procedure. It automatically formats and prints abend dumps. (See “Abend dump” on page 183.)

During ABDUMP processing, VTAM formats control blocks related to the abnormally ending task and prints them as part of the dump created by ABDUMP.

The following information shows ABDUMP formats the control blocks.

Note: This is an alphabetical list of the control blocks that might be in a dump. They might be in a different order in the dump.

Control block

Description

ACDEB

VTAM data extent block for the abnormally ending task

APPCB

LU 6.2 control block

COPR Control operator control block associated with the abnormally ending task

CRA Component recovery area for the abnormally ending task

FMCB Function management control block and extensions for the abnormally ending task

HSICB

Half-session information control block for the abnormally ending task

LUCB Logical unit control block associated with the abnormally ending task

MPST Memory-process scheduling table for the abnormally ending task

NSICB

Logical-network-services information control block for the abnormally ending task

NSSCB

Logical-network-services storage control block for the abnormally ending task

PST Process scheduling table for the abnormally ending task

RAB LU 6.2 resource allocation block for the abnormally ending task

RDTE Resource-definition-table application program entry for the abnormally ending task

SAB LU 6.2 logical-resource manager-session allocation block for the abnormally ending task

The following information shows formatted data areas described in :

ACDEB

CRA

FMCB

LUCB

MPST

PST

RDTE

The following information appears for each control block:

- A header line with the name and hexadecimal address of the beginning of the control block
- Under the header, the name of each selected field (as it appears in that control block's mapping DSECT) and the contents of the field (listed sequentially)
- After the formatted printout, a hexadecimal dump of the entire control block

SADMP service aid

SADMP formats and prints low-speed stand-alone dumps. During SADMP processing, VTAM formats selected control blocks and prints them as part of the dump created by SADMP.

See the diagnostic manuals for your operating system for more information on SADMP. See "Stand-alone dump" on page 184 for more information on this dump.

Chapter 6. Using VTAM dump analysis tools

This topic covers the following information:

- “Enhanced VTAM dump analysis tools”
- “Using VTAM interactive problem control system (IPCS) CLISTs” on page 192
- “Sample VTAM dump analysis functions” on page 193
- “VTAM formatted dump procedures” on page 196

Enhanced VTAM dump analysis tools

The VTAM dump analysis tools are enhancements to the IPCS subcommand VERBEXIT VTAMMAP. To use VERBEXIT VTAMMAP you can:

- Use the interactive panel interface.
- Enter VERBEXIT VTAMMAP subcommands from the IPCS command line.
- Create a batch job to issue the VERBEXIT VTAMMAP subcommands.

VTAMMAP will process SVC dumps, high-speed stand-alone dumps, or abend dumps.

“Sample VTAM dump analysis functions” on page 193 shows how you can access the VTAM dump analysis tools.

If you experience problems that you suspect to be related to the VTAM dump analysis tools, see “VTAM dump analysis tool problems” on page 47 for help.

Operating environment

The following rules apply:

- You cannot invoke multiple functions simultaneously using the IPCS command-line interface to formatted dump. Batch jobs do allow this with multiple calls to VTAMMAP, but only one command can be entered on a line.
- The VTAM-supplied IPCS CLISTs must be used with VTAM Version 4 Release 2 or higher.
- Multicultural support is not provided.
- IPCS is required for VTAM formatted dump.
- Many of the dump analysis tools described in this topic analyze control blocks that reside in VTAM private storage. If the tool cannot access VTAM private storage, the tool will not run correctly.
- The CLISTs described in this topic do *not* verify the accuracy of hexadecimal values such as storage addresses unless invoked from the ISPF panel.
- If you enter hexadecimal data either on the IPCS command line or using the batch option, you must enclose it in two sets of single quotation marks. (Do not use two sets if you are using the panel interface.)

IPCS strips off the first set of quotation marks, and the second set identifies hexadecimal data. For example, to be processed correctly, the string '02C72020' must be entered as "02C72020". The following example shows how to enter hexadecimal data in a command on the IPCS command line or in a batch job:

```
VERBEXIT VTAMMAP 'SIBCHECK ADDR(X' '01267B8' )'
```

- The VTAM internal trace (VIT) table has moved from ECSA storage to HVCOMMON storage in z/OS V1R13 Communications Server. To move the table, internal modifications were necessary, which resulted in an incompatibility with previous releases. You can still use the VTAM dump analysis tools VITAL, VTBASIC, VTVIT, and ALL on dumps from z/OS V1R12 Communications Server or earlier releases. However, these tools from z/OS V1R12 Communications Server or earlier releases will not operate correctly on dumps from z/OS V1R13 Communications Server or later releases.

Using VTAM interactive problem control system (IPCS) CLISTS

For the VTAM dump analysis functions that are not part of the enhanced dump analysis tools, you can use IPCS CLISTS to issue commands to analyze dumps of VTAM storage. To start the CLIST, type the CLIST name on the IPCS command line.

For example, to start ISTVABND, on the command line, type
ISTVABND

You can also use the CLIST command interface provided by IPCS to group TSO and IPCS commands together if you want to automate dump analysis procedures. See Table 48 on page 649 for a list of books that describe how to use IPCS.

Although these CLISTS are normally used online with IPCS, you can also issue them from the panel interface or run them in the background as batch jobs.

The IPCS CLISTS included in VTAM and described in this topic are:

- “ISTVABND” on page 242
- “ISTVDUMP” on page 244
- “ISTVMAP” on page 247
- “ISTVSAVE” on page 249
- “ISTVSLIP” on page 251

Obtaining online help for CLISTS

Note: This help operand applies only to the CLISTS; online help is invoked differently in the panel interface.

Each CLIST has online help information. To display it, enter the CLIST name followed by the HELP operand. Use no other operands, required or optional. For example, the following entry would display information on the ISTVDUMP CLIST.

```
ISTVDUMP HELP
```

After HELP information is displayed, you are prompted to either run the CLIST or exit the program. If you run the CLIST, you are prompted for each required operand.

Debugging CLIST errors

Note: This DEBUG option does not apply to the panel interface.

When you suspect an error in the execution of a CLIST, use the DEBUG option to list each command within the CLIST before and after execution. Specify the DEBUG option after any required parameters when the CLIST is invoked, as shown in the following example.

```
ISTVSLIP DEBUG
```

Printing CLIST output

The output from each CLIST is put into IPCSPRNT, the IPCS PRINT file. See Table 48 on page 649 to determine the document that contains information on IPCSPRNT.

Sample VTAM dump analysis functions

The following sample procedures provide examples of the ways in which you may access the VTAM dump analysis tools.

Using the panel interface

You can access the VTAM dump analysis tools by using the panel interface. These steps provide the minimum information that you need to use the panel interface.

Before you begin

You need to set the IPCS default dump to the data set name of the dump to be analyzed. You also need to set the IPCS options to direct the output either to print, or to the terminal, or both. See *z/OS MVS IPCS User's Guide* for more information.

Procedure

Perform the following steps to access VTAM formatted dump using the panel interface:

1. Log on to TSO

2. Access IPCS

3. Select option 7 from the option list.

```
-----IPCS PRIMARY OPTION MENU-----
OPTION  ===>  _

0  DEFAULTS  - Specify default dump and options
1  BROWSE   - Browse dump data set
2  ANALYSIS - Analyze dump contents
3  SUBMIT   - Submit problem analysis job to batch
4  COMMAND  - Enter IPCS subcommand or CLIST
5  UTILITY  - Perform utility functions
6  DUMPS    - Manage dump inventory
7  VTAM     - VTAM dump analysis
T  TUTORIAL - Learn how to use the IPCS dialog
X  EXIT     - Terminate using log and list defaults

Enter END command to terminate IPCS dialog
```

If you want a customized interface to be active to select VTAM, see z/OS Communications Server: New Function Summary for information on how to customize IPCS panel BLSPPRIM.

4. Select an option from the VTAMMAP Analysis Menu.

```
ISTD0001          VTAMMAP Analysis Menu

Select one of the following items. Then press Enter.

  1. APPC . . - APPLCONV, PARTNRLU, APPLMODE, APPMODAL
  2. APPN . . - APPNBASE, FNDADJCP, FNDANDCB, FND COS, FNDDECB, etc
  3. General. - HOST, VTAM, VTBASIC, VTFNDMOD, VTMODS, VITAL, etc
  4. Queues . - PABSCAN, VTCVTPAB, VTREADYQ
  5. Resource - RDTCHECK, RDTFULL, RDTHIER, RDTSUM, VTNODE
  6. Session. - ATMDATA, FINDDSIB, FINSIB, MNPS, SES, SIBCHECK
  7. Search . - SRTFIND
  8. Storage. - SPANC, STORAGE, VTBUF, VTRPH
  9. CSM . . - CSMALL, CSMBUF, CSMCMPID, CSMOWNER, CSMPOOL
 10. Waits. . - VTWRE
 11. ERs/VRs. - ROUTES, VTVRBLK
 12. CLISTs. - ISTVABND, ISTVDUMP, ISTVMAP, ISTVSAVE, ISTVSLIP
 13. APPN2. . - TRSTRACE

(C) Copyright IBM Corporation 1993,2006. All rights reserved.
Command ==>
F1=Help   F2=Split  F3=Exit   F9=Swap   F12=Cancel
```

5. Follow the screen prompts to process your dump.

Results

You know that you are done when a function is selected and the function executes. The output will go to the destination set in IPCS.

Using the IPCS command line

You can access the VTAM dump analysis tools by using the IPCS command line. These steps provide the minimum information that you need to use the IPCS command line.

Before you begin

You need to set the IPCS default dump to the data set name of the dump to be analyzed. You also need to set the IPCS options to direct the output either to print, or to the terminal, or both. See z/OS MVS IPCS User's Guide for more information.

Procedure

Perform the following steps to access VTAM formatted dump using the IPCS command-line interface.

1. Log on to TSO

2. Access IPCS

3. Select option 4 from the option list.

```

-----IPCS PRIMARY OPTION MENU-----
OPTION  ===>  _

0  DEFAULTS - Specify default dump and options
1  BROWSE   - Browse dump data set
2  ANALYSIS - Analyze dump contents
3  SUBMIT   - Submit problem analysis job to batch
4  COMMAND  - Enter IPCS subcommand or CLIST
5  UTILITY  - Perform utility functions
6  DUMPS    - Manage dump inventory
7  VTAM     - VTAM dump analysis
T  TUTORIAL - Learn how to use the IPCS dialog
X  EXIT     - Terminate using log and list defaults

Enter END command to terminate IPCS dialog

```

-
4. Enter a VTAMMAP command on the IPCS command line. For example:
`VERBEXIT VTAMMAP 'SIBCHECK ADDR(X''01267B8'')`
-

Results

You know that you are done when you type a command and the function executes. The output will go to the destination set in IPCS.

Using the batch option

You can access the VTAM dump analysis tools by using the batch option. These steps provide the minimum information that you need to use the batch option.

Before you begin

You need to find the data set name of the dump to be analyzed. (This name must be specified in the JCL.)

Procedure

Perform the following step to access VTAM formatted dump using the batch processing interface.

1. Prepare the JCL data set. See Table 48 on page 649 to determine what document describes IPCS.

Examples: Sample command (single command):

```
VERBEXIT VTAMMAP 'RDTFULL'
```

Sample command (multiple commands):

```
VERBEXIT VTAMMAP 'RDTFULL'
VERBEXIT VTAMMAP 'SIBCHECK ADDR(X''01267B8'')
```

-
2. Submit the JCL so the job will execute.
-

Results

You know that you are done when the job completes.

VTAM formatted dump procedures

This topic contains an alphabetical list of the VTAM formatted dump analysis tools and IPCS CLISTs available with VTAM.

The descriptions for each tool include:

- Procedure name
- Description
- Operands
- Syntax
- Sample output

ALL

Use ALL to invoke the following functions:

RDTFULL
ROUTES
SES
STORAGE
VTAM
VTBASIC

Note: The ALL function is not displayed on the main formatted dump panel with other General functions, but is available on the General panel.

Operands

Trace output

Enter **Format** to format the VIT and **No format** to display the VIT in hexadecimal format. **Format** is the default.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

ALL

ALL Analysis

VTAM INTERNAL TRACE TABLE 000001EF_81000000

PRESENT WRAP C728E7D3 D61C13C4 LAST WRAP 00000000 00000000

CURRENT ENTRY 000001EF_83503020 LAST ENTRY 000001EF_841FFFE0

C4E2D740 12582410 02915E88 00CC4908 02A275F8 02A275F8 E3E2E6E4 02929010

D3D2E2C8 12000100 00CC4C70 00000000 82A95442 00000000 00000000 02929010

E4D5D3D2 12000100 00CC4C70 00000100 82A9546C 00000000 01000000 02929010

D8E4C558 12482810 02915E88 00CC4248 82A954F8 02A275F8 C9D5E3D4 02929010


```

:
:
ATCVT: 00CC41F8
  ATCRDT... 02955740  ATCSRT... 02C35008  ATCCONFT. 00CC18E8
  ATCBPDA.. 02953000  ATCACTRM. 0000      ATCVTL0D. 02A27650
  GWSSCP = YES
DATA: 00CC41F8
+0000 E5C5F4F3  40404040  FFF900C8  02825000  | VE43   .9.H.b&. |
+0010 00000000  0000FFF9  11280000  00000000  | .....9..... |
+0020 02915E88  00000000  00000000  00000000  | .j;h..... |
+0030 00CC4524  00000000  13201000  00000010  | ."..... |
+0040 11280000  00000000  02915E88  00000000  | .....j;h.... |
:
:
RDTE: 02CE0CEC
  RPRNAME.. APPCAP09  RPRENTRY. 55      RPRBITAN. 09000810  01
  RPRDEVCH. C06D0000  00800000
DATA: 02CE0CEC
+0000 C1D7D7C3  C1D7F0F9  80000000  00550200  | APPCAP09..... |
+0010 00000000  00010095  000A0000  02CE0DE8  | .....n.....".Y |
+0030 02CE0008  00000000  00000000  02000200  | ."..... |
+0040 00090008  10010010  00000000  00000000  | ..... |
:
:
No SIBs on the ATCSIBQ chain

```

APPLCONV

Use APPLCONV to display all conversations for an APPC application. APPLCONV formats and displays the APPCB control block, and the COPR control block if present. It will also format and display the APPC resource allocation block (RAB) and each session control block (SAB) associated with the RAB.

Operands

APPC application name

The APPC application name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

The APPC application name is required.

Syntax

```

▶▶—APPLCONV— —APPLNAME—(—APPC_application_name—)————▶▶

```

Sample output

APPLCONV APPLNAME(APPCAP05)

```

                                APPLCONV Analysis
APPCB: 0290E6B8
  APPLUCB.. 0291A100  APPTSKID. 02818588  APPACB... 00CB4820
  APPLUN... APPCAP05
  APPSPTAE 02906530 0290D898 02906620 029065D0 02906580
DATA: 0290E6B8
+0000 62C1D7D7  0291A100  02818588  00000000  | .APP.j~.aeh.... |
+0010 00000000  00000000  024100B0  00000000  | .....[.... |
+0020 31094000  00000010  00000000  00000000  | ..... |
+0030 0101001B  00000000  00CB4820  0290B088  | .....[h |
+0040 00000000  00000000  00000000  00000000  | ..... |
:
:
No session limit negotiations were in progress

                                Current conversation(s) for APPCAP05
RAB: 02804028

```

APPLCONV

```
RABCONID. 01000003 RABCRPLA. 00000000 RABPSFSM. 01000000 04800000
RABSABPT. 028030C0 RABNETID. NETA RABLUNAM. APPCAP06
RABMODEN. BATCH
Conversation State SEND
DATA: 02804028
+0000 62D9C1C2 00000000 02818588 00000000 |.RAB.....aeh....|
+0010 00000000 01000003 01000004 D5C5E3C1 |.....NETA|
+0020 40404040 C1D7D7C3 C1D7F0F6 C2C1E3C3 | APPCAP06BATIC|
+0030 C8404040 A50F95D0 028030C0 00000000 | H v.n}...{....|
+0040 00000000 00000000 024100B4 00000000 |.....+....|
:
SAB: 028030C0
SABSHARE. C0 SABLRFML. 00 SABFSM... 30
SABSSENSE. 00000000 SABNSFG1. 03 SABNSFG2. 60
DATA: 028030C0
+0000 62E2C1C2 00000000 00000000 01000004 |.SAB.....|
+0010 02804028 D5C5E3C1 40404040 C1D7D7C3 |..NETA APPC|
+0020 C1D7F0F6 C2C1E3C3 C8404040 C0003000 |AP06BATIC {...|
+0030 00000000 00000000 02803028 00000000 |.....|
+0040 00000000 00ABEEC3 CE09D9CC 00000000 |.....z.C0.R"|
+0050 08000360 00000000 00000000 00000000 |...-.....|
+0060 00000000 00000000 00000000 00000000 |.....|
+0070 00000000 00000000 00000000 00000000 |.....|
+0080 00000000 00000000 00000000 |.....|
```

APPLMODE

Use APPLMODE to display all logon modes in the logon mode table for conversations between an application program and a particular partner LU. APPLMODE will process the LU entries searching for all LU entries that match the specified partner LU name and the optional partner LU network identifier. If entries are found, APPLMODE processes the chain of modes and determines the settings for:

- Current session limits(x,y,z), where:
 - x - Session limit
 - y - Minimum number of contention winner sessions for local LU
 - z - Minimum number of contention winner sessions for remote LU
- Current session count(x,y,z), where:
 - x - Active session count
 - y - Active contention winners at local LU
 - z - Active contention winners at remote LU
- Pending session counts(x,y,z), where:
 - x - Count of pending sessions
 - y - Count of pending contention winners
 - z - Count of pending contention losers
- Pending session termination counts(x,y,z), where:
 - x - Pending termination contention winners
 - + Pending termination contention losers
 - y - Pending termination contention winners
 - z - Pending termination contention losers
- Defined session limits(x,y,z), where:
 - x - Defined session limit
 - y - Defined minimum number of contention winner sessions for local LU
 - z - Defined minimum number of contention winner sessions for remote LU

APPLMODE formats and displays the APPCB control block, and the COPR control block if present.

Operands

APPC application name

The APPC application name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

The APPC application name is required.

Partner LU name

The partner LU name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

The partner LU name is required.

Partner LU NetID

The partner LU NetID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

Use the following syntax as an alternative to the panel interface.

Syntax

```

▶▶—APPLMODE— —APPLNAME—(—APPC_application_name—)—————▶
▶ —LU—(—partner_LU_name—)—————▶
      └── —NETID—(—partner_LU_netid—) ───┘

```

Sample output

APPLMODE APPLNAME(APPCAP05) LU(APPCAP06)

```

                                APPLMODE Analysis
APPCB: 0290E6B8
APPLUCB.. 0291A100  APPTSKID. 02818588  APPACB... 00CB4820
APPLUN...  APCCAP05
APPSPTAE 02906530 0290D898 02906620 029065D0 02906580
DATA: 0290E6B8
+0000 62C1D7D7 0291A100 02818588 00000000 |.APP.j~..aeh....
+0010 00000000 00000000 024100B0 00000000 |.....[....
+0020 31094000 00000010 00000000 00000000 |.....
+0030 0101001B 00000000 00CB4820 0290B088 |.....[h
+0040 00000000 00000000 00000000 00000000 |.....
+0050 00000000 00000000 00000000 00000000 |.....
+0060 C1D7D7C3 C1D7F0F5 00000000 00000000 |APCCAP05.....
+0070 00000000 00000000 028030C0 02804028 |.....{...
+0080 00000000 0292A2FC 00000000 00000000 |.....ks.....
+0090 00000000 0290DA98 00000000 00000000 |.....q.....
+00A0 00000000 00000000 023FA0A8 00000000 |.....μy.....
+00B0 36200000 000000A0 02906530 0290D898 |.....μ.....Qq
+00C0 02906620 029065D0 02906580 00000000 |.....}.....
+00D0 00000000 00000000 00000000 00000000 |.....
+00E0 00000000 00000000 00000000 00000000 |.....
+00F0 00000000 00000000 00000000 00000000 |.....
+0100 00000000 00000000 |.....

```

No session limit negotiations were in progress

Modes between application APCCAP05 and partner LU APCCAP06

Mode name SNASVCMG

```

Current session limits      (X'0002',X'0001',X'0001')
Current session counts     (X'0001',X'0001',X'0000')
Pending session counts     (X'0000',X'0000',X'0000')

```

APPLMODE

```
Pending session termination counts (X'00000000',X'0000',X'0000')
Define session counts              (X'0004',X'0002',X'0002')

Mode name BATCH

Current session limits             (X'0004',X'0002',X'0002')
Current session counts             (X'0001',X'0001',X'0000')
Pending session counts             (X'0000',X'0000',X'0000')
Pending session termination counts (X'00000000',X'0000',X'0000')
Define session counts              (X'0004',X'0002',X'0002')
```

APPMODAL

Use APPMODAL to display all information about a particular logon mode for a conversation between an application and a partner LU. APPMODAL will process the LU entries searching for all LU entries that match the specified partner LU name and the optional partner LU NetID. If LU entries are found, the chain of modes is searched for a match to the specified logon mode name. If a matching logon mode is found, APPMODAL determines the settings for:

- Current session limits(x,y,z), where:
 - x - Session limit
 - y - Minimum number of contention winner sessions for local LU
 - z - Minimum number of contention winner sessions for remote LU
- Current session count(x,y,z), where:
 - x - Active session count
 - y - Active contention winners at local LU
 - z - Active contention winners at remote LU
- Pending session counts(x,y,z), where:
 - x - Count of pending sessions
 - y - Count of pending contention winners
 - z - Count of pending contention losers
- Pending session termination counts(x,y,z), where:
 - x - Pending termination contention winners + losers
 - y - Pending termination contention winners
 - z - Pending termination contention losers
- Defined session limits(x,y,z), where:
 - x - Defined session limit
 - y - Defined minimum number of contention winner sessions for local LU
 - z - Defined minimum number of contention winner sessions for remote LU

APPMODAL displays:

- Active conversations between the two applications on the logon mode by running the chain of RABs
- Waiting requests off the logon mode (requests for conversations that have not been serviced)
- Free sessions on the logon mode (SABs that represent sessions that are not currently assigned to a conversation)
- Pending active sessions on the logon mode (SABs that represent sessions that are in the process of being activated on the logon mode)

APPMODAL formats and displays the APPCB, COPR, LME (selected fields), RAB, and SAB control blocks.

Operands

APPC application name

The APPC application name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

The APPC application name is required.

Partner LU name

The partner LU name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

The partner LU name is required.

Partner LU NetID

The partner LU NetID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.


Logon mode name

The logon mode name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

The logon mode name is required.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

APPMODAL APPLNAME(APPCAP05) LU(APPCAP06) LOGMODE(BATCH)

```

                                APPMODAL Analysis

APPCB: 0290E6B8
  APPLUCB.. 0291A100  APPTSKID. 02818588  APPACB... 00CB4820
  APPLUN... APPCAP05
  APPSPTAE 02906530 0290D898 02906620 029065D0 02906580
DATA: 0290E6B8
+0000 62C1D7D7 0291A100 02818588 00000000 | .APP.j~..aeh.... |
+0010 00000000 00000000 024100B0 00000000 | .....[.... |
+0020 31094000 00000010 00000000 00000000 | .. ..... |
+0030 0101001B 00000000 00CB4820 0290B088 | .....[h |
+0040 00000000 00000000 00000000 00000000 | ..... |
:
No session limit negotiations were in progress

LME: 0290B148
  LMENETID. NETA      LMENM.... APPCAP06  LMEFSM... C2

Mode name BATCH

Current session limits      (X'0004',X'0002',X'0002')
Current session counts     (X'0001',X'0001',X'0000')
Pending session counts     (X'0000',X'0000',X'0000')
Pending session termination counts (X'00000000',X'0000',X'0000')
Define session counts      (X'0004',X'0002',X'0002')
  
```

APPMODAL

Current conversation(s)

```
RAB: 02804028
  RABCONID. 01000003  RABCRPLA. 00000000  RABPSFSM. 01000000  04800000
  RABSABPT. 028030C0  RABNETID. NETA      RABLUNAM. APPCAP06
  RABMODEN. BATCH
  Conversation State SEND
DATA: 02804028
+0000 62D9C1C2 00000000 02818588 00000000 | .RAB.....aeh.... |
+0010 00000000 01000003 01000004 D5C5E3C1 | .....NETA       |
+0020 40404040 C1D7D7C3 C1D7F0F6 C2C1E3C3 |      APPCAP06BATC |
+0030 C8404040 A50F95D0 028030C0 00000000 | H  v.n}...{.... |
+0040 00000000 00000000 024100B4 00000000 | .....+..... |
:
:
No conversations found awaiting BID response

No free sessions found

No pending active sessions found
```

APPNBASE

Use APPNBASE to format the global APPN control blocks:

- ACMDT
- APNVT
- DRDAT
- MTDAT
- SCDAT
- SLGDT
- TRDAT

The control block addresses and the hexadecimal data from each control block are provided to help you diagnose APPN problems.

Use the following syntax as an alternative to the panel interface.

Syntax

▶▶ APPNBASE ◀◀

Sample output

APPNBASE

APPNBASE Analysis

```
APNVT: 00C1BD40
+0000 C1D7D7D5 00000000 062FEE88 00C1C118 | APPN.....h.AA. |
+0010 0652A948 00000000 00000000 00000000 | ..Z..... |
+0020 00000000 00000000 00000000 00000000 | ..... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 068D8E08 00000000 | ..... |
+0050 062EBE68 00000000 00C1BDC4 00000000 | .._.....A D.... |
:
:
ACMDT: 00C1C118
+0000 C1C3D4C4 0004C1F0 F1D50000 00000004 | ACMD..A01N..... |
+0010 D5C5E3C1 00000000 09D5C5E3 C14BC1F0 | NETA.....NETA.A0 |
+0020 F1D54040 40404040 40408080 00000000 | 1N          ..... |
```

```

+0030 00000000 00000000 80380100 00000000 | ..... |
+0040 000A8C00 00000000 00000024 00006C60 | .....%- |
+0050 00000064 00000000 00000000 00000000 | ..... |
:
DRDAT: 062EB1A0
+0000 C4D9C4E3 0000C800 C0000018 069EF200 | DRDT..H.{.....2. |
+0010 069EF110 069EF2F0 00041000 C0000018 | ..1...20....{... |
+0020 069EF098 069EF098 069EF098 00041000 | ..0q..0q..0q... |
+0030 8000002C 069EF110 00480400 00000000 | .....1..... |
+0040 00000000 00000000 00000000 00000000 | ..... |
+0050 00000000 824B5D76 8686D5D0 824B4E9A | ....b.) .ffN}b.+ |
:
MTDAT: 062EBD70
+0000 D4E3C4E3 00000000 00000000 00000000 | MTDT..... |
+0010 00000000 068DFFA8 068DFF48 067EDEF8 | .....y_...=.8 |
+0020 067EDE70 067EFFD0 067EFF98 068DE008 | .=...=.}.=.q_\. |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 0000003C 00000000 | ..... |
+0050 10300000 00000000 00000000 80000004 | ..... |
:
SCDAT: 062D1B88
+0000 E2C3C4E3 00000000 C0000088 069C6B10 | SCDT....{..h... |
+0010 069C6138 069C67C8 00041100 00000000 | ../.H..... |
+0020 00000000 00000000 00000000 00000000 | ..... |
+0030 068E9C18 00000000 00000000 00000000 | ..... |
+0040 00000000 069C6480 60C3D7E2 E5C3D4C7 | .....-CPSVCMG |
+0050 40000000 00000000 000C12C1 00000000 | .....A.... |
:
SLGDT: 062EB080
+0000 E2D3C4E3 00000000 8000002C 06CCBA20 | SLDT..... |
+0010 00082400 00000000 00000000 40000024 | ..... |
+0020 06ACBAA0 06C10020 40000004 00000000 | .....A. .... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 824B5D76 868BDC10 824B4E9A 06318010 | b.) .f...b.+... |
+0050 00C1BED0 06B6A100 062FEE88 00000C60 | .A_}.....h...- |
:
TRDAT: 062D1848
+0000 E3D9C4E3 068F0008 067ACF90 40000000 | TRDT.....:.. |
+0010 00000000 00000000 40000000 00000000 | ..... |
+0020 00000000 00000000 4000001C 00000000 | ..... |
+0030 00000000 40000008 00000000 00000000 | ..... |
+0040 10000000 00000000 40000008 00000000 | ..... |
+0050 00000000 00000000 00000000 00000000 | ..... |
:

```

ATMDATA

Use ATMDATA to format control blocks associated with ATM support.

Operands

Line name

The line name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

Major node name

The major node name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

Port name


The port name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

ATMDATA

You can specify one or more operands. If you specify more than one operand, the associations among the operands must be correct, or you will not get any control blocks. For example, if you specify line name and major node name, the line name must be the name of a LINE definition statement in the major node you specify for major node name.

You can also specify ATMDATA without any of the operands. If you do, you will get the control blocks associated with all of the lines and ports defined in all of the XCA major nodes.

Syntax



Sample output

ATMDATA PORTN(OSA2ATM1)

ATMDATA Analysis

```

ALPOR: 08374018
+0000 C1D3D7D6 00000000 0847F048 00000000 | ALPO.....0.....
+0010 00000000 00000000 0839C04C 00000000 | .....{<....
+0020 B2090000 0FF00010 00000000 00000000 | .....0.....
+0030 08000000 00000000 08373018 083740B0 | .....
+0040 00000000 00000000 470E470E 63B70000 | .....
+0050 D6E2C1F2 C1E3D4F1 00010102 88C923E4 | OSA2ATM1...hI.U
+0060 88C923E4 88C923E4 88C923E4 00000000 | hI.UhI.UhI.U...
+0070 00020014 39999999 99999999 99999801 | .....rrrrrrrrr.
+0080 01315045 53543407 01000000 00000000 | ..&.....
+0090 E7C3C1D6 E2C1F240 C1D3C6C9 | XCAOSA2 ALFI
  
```

```

ALFIL: 083740B0
+0000 C1D3C6C9 01000000 00000000 00000000 | ALFI.....
+0010 00000000 | ....
  
```

```

OSLIN: 07F743F0
+0000 D6000000 01000000 C0010000 00000000 | O.....{.....
+0010 08372018 00000000 00000000 00000000 | .....
+0020 00000000 00000000 00000000 00000000 | .....
+0030 00000000 00000000 00000000 00000000 | .....
+0040 00000000 00000000 00000000 00000000 | .....
  
```

```

ALNCB: 08372018
+0000 2FD000E8 08FE83B8 0847F048 00050000 | }.Y..c...0....
+0010 00000004 00000000 00000000 00000000 | .....
+0020 00000000 00000000 00000000 0900A010 | .....
+0030 00000000 00000000 00000000 00000000 | .....
+0040 00000000 00000000 0839C06C 00000000 | .....{%.
+0050 B1090000 0FF00040 00000000 00000000 | .....0. ....
+0060 00000000 00000000 00000000 00000000 | .....
+0070 00000000 00000000 00000000 00000000 | .....
+0080 00000000 00000000 0839C048 00000000 | .....{....
+0090 B0090000 1FF00080 00000000 00000000 | .....0. ....
+00A0 00000000 00000000 00000000 00000000 | .....
+00B0 00000000 00000000 00000000 00000000 | .....
+00C0 00000000 00000000 00000000 00000000 | .....
+00D0 00000000 00000000 00000000 00000000 | .....
  
```


+00E0	00000000	00000000	00000000	00000000
+00F0	00000000	3000C000	00000000	0837210C{.....
+0100	00000000	00000000	08368080	00000000
+0110	00000000	00000000	00000000	00000000
+0120	00000000	00000000	00000000	00000000
+0130	00000000	00000000	00000000	00000000
+0140	00000000	00000000	00000000	00000000
+0150	00000000	00000000	00000000	00000000
+0160	08353018	00000000	00000000	00000000
+0170	00000000	00000000	00000000	00000000
+0180	1B100000	08372018	00000000	00000000
+0190	58588000	00000000	00000000	00000000
+01A0	00000000	05000000	00000000	00000000
+01B0	00000000	00000000	00000000	00000000
+01C0	00000000	00000000	00000000	00000000
+01D0	00000000	00000000	00000000	00000000
+01E0	ACA9B06C	E45A8806	00000000	00000000	.z.%U!h.....
+01F0	00000000	08374018	88C923E4	88C923E4hI.UhI.U
+0200	88C923E4	08371018	00000000	00000000	hI.U.....
+0210	88C7E014	08371810	00000000	00000000	hG\.....
+0220	00000000	0006000D	000A0501	0200470E
+0230	00000000	00C0470E	470E470E	00E90000{.....Z..
+0240	C8000002	00143999	99999999	99999999	H.....rrrrrrrrr
+0250	98010131	50455354	34070000	00000000	q...&.....
+0260	00000000	00000000	00000000	00000000
+0270	00000000	00000000	00000000	00000000
+0280	00000000	00000000	00000000	00000000
+0290	00000000	00000000	00000000	00000000
+02A0	00000000	00010102	00010108	00010108
+02B0	D6E2C1F2	D7E5C3F1	00000004	01000101	OSA2PVC1.....
+02C0	00000000	FF808021	01000000	00

CSMALL

CSMALL displays the major control block, a summary of all CSM pools, and each pool control block and its associated extent control blocks. In the following pool summary, only the pools which have been created are displayed.

Use the following syntax as an alternative to the panel interface.

Syntax

▶▶—CSMALL—▶▶

Equated symbol

Symbol	Description
--------	-------------

IVTDSpace

The last CSM data space buffer processed

Sample output

CSMALL

CSMALL Analysis

IVTSMCST:	03E1B000					
+0000	C3E2D440	434C0000	83D75A78	83D7CE58	CSM .<..cP!.cP..	
+0010	03EC1120	00640000	000067AC	00640000	
+0020	0000643D	83D7B0A8	83D76DC0	83D82AB8cP.ycP_.cQ..	

CSMALL

+0030	00800000	03EC1198	00000000	00025F88q.....h
+0040	00000000	0002A190	03EC1128	03ED8000
+0050	03E20000	00000039	00000000	00000002	.S.....
+0060	00000002	00000018	F0000800	032F80000.....
+0070	80000000	00000000	01C432F8	03338000D.8....
+0080	F0000800	02AF3000	00000000	00000000	0.....
+0090	01C42AF3	02B33000	00000000	00000000	.D.3.....
+00A0	00000000	00000000	00000000	00000000
+00B0	00000000	00000000	00000000	00000000
+00C0	00000000	00000000	00000000	00000000
+00D0	00000000	00000000	00000000	00000000
+00E0	C3E2D46D	C7D9E26D	D3C1E3C3	C8E2C5E3
+00F0	40404040	40404040	40404040	40404040	CSM_GRS_LATCHSET
+0100	40404040	40404040	40404040	40404040
+0110	00E11038	C6559180	00000010	83D85D58F.j.....cQ).
+0120	83D86568	83D7D1E0	83D78278	83D79930	cQ...cPJ.cPb.cPr.
+0130	83D85AE8	83D73348	00000000	83ECD898	cQ!YcP.....c.Qq
+0140	83D83658	03EC1080	03EC10F0	03EC1104	cQ.....0....
+0150	03EC1106	83ECC5B0	83ED6F78	03B47D68c.E.c?...!
+0160	83ECD520	03EC36C8	00000000	00000000	c.N....H.....
+0170	00000000	83ECC5B0	80EC1120	00000000c.E.....
+0180	00000000	03EC3768	00000000	000067A0
+0190	0000643D	00E065AC	00000000	00000000
+01A0	00E05830	00000000	E5C5F6F1	F2404040VE612
+01B0	F0F6F1F2	F5F6F9F5	60F1F1F7	F0F160F1	06125695-11701-1
+01C0	F2F00000	00000000	80E05730	00000000	20.....
+01D0	00000000	03D731DC	0000001C	83D86DE0P.....cQ..
+01E0	03D716F8	83D7F6F8	83D7CA28	83D80290	.P.8cP68cP..cQ..
+01F0	83D7A390	83D82518	83D85868	83D854D0	cPt.cQ..cQ..cQ..
+0200	83D82D80	83D84D68	83D84BE8	83D848F8	cQ..cQ(.cQ.YcQ.8
+0210	00550000	005A0000	00000B00	00000000!
+0220	0000000B	00000000	03EC3DF8	000000008....
+0230	00000000	00550000	005A0000	00001B00!
+0240	00004500	00000000	0000001B	00000000
+0250	00000000	00000000	00000000	00000000
+0260	00000000	00000000	03E18C00	00000008
+0270	07333333	00000000	03B45000	03B01000&.....
+0280	03AFE000	03AFB000	00000000	032F3000
+0290	02AEB000	02AE5000	02ADF000	02ADC000&..0.....
+02A0	03AF8000	02AEE000	02AE8000	02AE2000
+02B0	00000000	20000000	00000000	00000000
+02C0	00000000	00000000	00000000	00000000
+02D0	00000000	00000000	00000000	00000000
+02E0	00000000	00000000	00000000	00000000
+02F0	00000000	00000000	00000000	00000000
+0300	00000000	00000000	00000000	00000000
+0310	00000000	00000000	00000000	00000000
+0320	00000000	00000000	00000000	00000000
+0330	00000000	00000000	00000000	00000000
+0340	00000000	00000000	00000000	00000000

CSM Pool Summary

CSM	4K	ECSA	POOL	03B45000
CSM	16K	ECSA	POOL	03B01000
CSM	32K	ECSA	POOL	03AFE000
CSM	60K	ECSA	POOL	03AFB000
CSM	180K	ECSA	POOL	00000000
CSM	4K	DSPACE31	POOL	032F3000
CSM	16K	DSPACE31	POOL	02AEB000
CSM	32K	DSPACE31	POOL	02AE5000
CSM	60K	DSPACE31	POOL	02ADF000
CSM	180K	DSPACE31	POOL	02ADC000
CSM	4K	DSPACE64	POOL	03AF8000
CSM	16K	DSPACE64	POOL	02AEE000
CSM	32K	DSPACE64	POOL	02AE8000
CSM	60K	DSPACE64	POOL	02AE2000

CSM 180K DSPACE64 POOL 00000000

CSM 4K ECSA POOL

CSMPPOOL: 03B45000

+0000	D7D6D6D3	20608000	00001000	00000002	POOL.-.....
+0010	00000040	00000035	00000008	00000010
+0020	03E20630	03B14A80	00000002	02196010	.S.....-
+0030	00000000	00000040	00000000	00000000
+0040	00000000	03E1B278	03EC3D58	00000100
+0050	00000040	01100001	03E20630	00000000S.....

CSMEXT: 03E20630

+0000	C5E7E340	05800000	03E1AA80	00000000	EXT
+0010	03B45000	00000000	00000000	00000000	..&.....
+0020	00000000	03B35000	03B44FFF	00000010&.....
+0030	00000005	00000000	00010000	00000010
+0040	00000000	00000000	00000000	00000000
+0050	00000000	00000000	00000000	FFE00000
+0060	00000000	00000000	00000000	00000000
+0070	00000000	00000000	00000000	00000000

⋮

CSM 16K ECSA POOL

CSMPPOOL: 03B01000

+0000	D7D6D6D3	20608000	00004000	00000001	POOL.-....
+0010	00000000	00000000	00000004	00000000
+0020	00000000	00000000	00000001	03EDA020
+0030	00000000	00000004	00000000	00000000
+0040	00000000	03E1B27C	03EC31D8	00000400@...Q...
+0050	00000000	02200004	00000000	00000000

⋮

CSM 4K DSPACE31 POOL

CSMPPOOL: 032F3000

+0000	D7D6D6D3	20604000	00001000	00000001	POOL.-
+0010	00000040	00000040	00000008	00000010
+0020	032F6500	02AF1A80	00000001	03ED0000
+0030	00000000	00000040	00000000	00000000
+0040	00000000	03E1B28C	03EC0000	00000100
+0050	00000040	06900001	032F6500	00000000

CSMEXT: 032F6500

+0000	C5E7E340	05800000	02AF2A80	00000000	EXT
+0010	032F3000	01C42AF3	F0000800	02AF3000D.30.....
+0020	00000002	02AF3000	02B02FFF	00000010
+0030	00000010	00000000	00010000	00000010
+0040	F0C1C1C1	C3C3E2D4	00000000	00000000	CSM31002.....
+0050	00000000	00000000	00000000	00000000
+0060	00000000	00000000	00000000	00000000
+0070	00000000	00000000	00000000	00000000

⋮

CSM 16K DSPACE31 POOL

CSMPPOOL: 02AEB000

+0000	D7D6D6D3	20604000	00004000	00000001	POOL.- ...
+0010	00000000	00000000	00000004	00000000
+0020	00000000	00000000	00000001	03ECA010
+0030	00000000	00000004	00000000	00000000

- ASN2
- CHG2
- CPY3
- CPY4
- FIX2
- FRB2
- GTB3
- PAG2

Use the following syntax as an alternative to the panel interface.

Syntax

▶▶—CSMBUF— —CSMTOKEN—(—*buffer_token*—)————▶▶

Equated symbol

Symbol	Description
--------	-------------

IVTDSpace

The last CSM data space buffer processed

Sample output

CSMBUF CSMTOKEN(03E2063003E206B000000000)

CSMBUF Analysis

CSM	4K	ECSA	POOL		
CSMPOOL: 03B45000					
+0000	D7D6D6D3	20608000	00001000	00000002	POOL.-.....
+0010	00000040	00000035	00000008	00000010
+0020	03E20630	03B14A80	00000002	02196010	.S.....-
+0030	00000000	00000040	00000000	00000000
+0040	00000000	03E1B278	03EC3D58	00000100
+0050	00000040	01100001	03E20630	00000000S.....
CSMEXT: 03E20630					
+0000	C5E7E340	05800000	03E1AA80	00000000	EXT
+0010	03B45000	00000000	00000000	00000000	..&.....
+0020	00000000	03B35000	03B44FFF	00000010&.....
+0030	00000005	00000000	00010000	00000010
+0040	00000000	00000000	00000000	00000000
+0050	00000000	00000000	00000000	FFE00000
+0060	00000000	00000000	00000000	00000000
+0070	00000000	00000000	00000000	00000000
CSMHDR: 03E206B0					
+0000	C8C4D940	00502080	00020000	00000000	HDR .&.....
+0010	00000000	03B35000	03E206B0	00000000&..S.....
+0020	00000000	00000000	00000000	00000000
+0030	B5568781	8255488A	00000001	00000001	..gab.....
+0040	00000000	00000000	01000100	00000000
DATA: 03B35000					
+0000	00000000	00E00000	00000015	00000014
+0010	00000080	0FFC0001	00000000	81010001a...

CSMBUF

```

+0020 00000000 00000000 00240040 00004005 | ..... |
+0030 0001010A 00000000 00000000 01000040 | ..... |
+0040 00000040 C608D400 00000000 0000FF00 | ... F.M..... |
+0050 168765E5 00000013 0004000D 00000000 | .g.V..... |
+0060 00000173 03225902 00000000 00000000 | ..... |
+0070 050E0000 00010001 000000CC 00000000 | ..... |
+0080 00000000 D5C5E3C1 4BE2E2C3 D7F1C100 | ... NETA.SSCP1A. |
+0090 00000000 00000000 6012C20A 80000000 | .....-B..... |
+00A0 03000000 0010440B D5C5E3C1 4BE2E2C3 | .....NETA.SSC |
+00B0 D7F1C100 000C450A 80000000 0201083B | P1A..... |
+00C0 18054803 800F1646 1480150B D5C5E3C1 | .....NETA |
+00D0 4BE2E2C3 D7F2C121 80000001 16470000 | .SSCP2A..... |
+00E0 0002888D 00000000 00000000 014C0080 | ..h.....<.. |
+00F0 80800548 03800F01 00000000 81010001 | .....a... |
+0100 00000000 00000000 00240052 00005205 | ..... |
+0110 0001010A 00000000 00000000 01000052 | ..... |
+0120 00000120 C608D400 00000000 0000FF00 | ... F.M..... |
+0130 168765E4 00000011 3C04000D 00000012 | .g.U..... |
+0140 000000B9 03228520 00047520 00000000 | .....e..... |
+0150 050E0000 00010001 00000173 00000000 | ..... |
+0160 00000000 5D000000 00000000 00000001 | ...). |
+0170 83010000 00070000 00000000 00000000 | c..... |
+0180 00000000 00000000 00000000 00000000 | ..... |
:
+0FF0 00000000 00000000 00000000 00000000 | ..... |

```

CSMCMPID

Use CSMCMPID to display the addresses of all CSM buffers currently used by a specific component ID or by all component IDs.

This command summarizes each buffer currently owned by the specified component ID for each pool size and type combination and lists the following information:

- The pool
- One line for each extent (with data space name)
- One line for each buffer (both the primary and image headers are searched)
- The total number of buffers in each pool with that component ID

Only the pools that have been created are displayed.

Note: This command does not use the old component IDs of the buffers to list them.

In the extent line, if the buffers are in a data space pool, the data space name is shown in the **DSPNAME** field. If the buffers reside in ECSA, then the **DSPNAME** field is labeled N/A.

In the buffer line, if the buffer address is in an image header, the primary header address is shown in the **PHDR** field.

Operands

CSMCOMP(*component_identifier*)

Specifies the component identifier (compid). The component ID must be two hexadecimal characters in the form X'nn'. If not specified, defaults to all component IDs.

CSMOWNID(*owner_identifier*)

Specifies the owner identifier of the buffer. The owner ID is the address space


identifier (ASID). The owner ID must be two to four hexadecimal characters in the form X'nn'. If less than four characters, the owner ID is padded on the left with zeros.

ALLCMPID

When ALLCMPID is specified with the DETAIL operand, all component IDs with buffers are displayed. When ALLCMPID is specified without the DETAIL operand, the summary of all component IDs with the buffers is displayed for all CSM pools.

The default is ALLCMPID without DETAIL.

Syntax



Equated symbol

Symbol

Description

IVTDSpace

The last CSM data space buffer processed

Sample output

CSMCMPID CSMCOMP(01)

CSMCMPID Analysis

CSM 4K ECSA Pool 0CD03000 for Component ID 01

```

EXT 0DC6C4B8 DSPNAME N/A
HDR 0DC6C588 PHDR N/A      COMPID 00000001 USER 00000000 BUF 0CCF4000
HDR 0DC6C628 PHDR N/A      COMPID 00000001 USER 00000000 BUF 0CCF5000
HDR 0DC6C678 PHDR N/A      COMPID 00000001 USER 00000000 BUF 0CCF6000
"
"
"
  
```

```

Number of primary buffers owned by compid 03 in 4K ECSA pool: 263
Number of _image_ buffers owned by compid 02 in 4K ECSA pool: 84
  
```

This output is repeated for all buffer pools with buffers.

CSMCMPID ALLCMPID

CSMCMPID Analysis

CSM ALL COMPID Summary

```

CSM 4K ECSA Pool 2B4E6000 for Component ID ALL
Number of primary buffers owned by COMPID 20 in 4K ECSA pool: 1
Number of _image_ buffers owned by COMPID 20 in 4K ECSA pool: 0

Number of primary buffers owned by COMPID 94 in 4K ECSA pool: 15
Number of _image_ buffers owned by COMPID 94 in 4K ECSA pool: 51
  
```

CSMCOMPID

```
Number of primary buffers owned by COMPID B1 in 4K ECSA pool: 121
Number of _image_ buffers owned by COMPID B1 in 4K ECSA pool: 560
"
"
"
```

This output is repeated for all buffer pools with buffers.

CSMOWNER

Use CSMOWNER to display the addresses of all CSM buffers owned by a specific user.

This command summarizes each buffer owned by the specified owner ID for each pool size and type combination. First, the pool is listed, then one line for each buffer (both the primary and image headers are searched), then a total showing how many buffers the owner owns in each pool. Only the pools which have been created are displayed.

In the buffer line, if the buffer address is in an image header, the primary header address is shown in the **PHDR** field. Otherwise, the **PHDR** field is labeled N/A. If the buffer is in a data space pool, the data space name is shown in the **DSPNAME** field. If the buffer resides in ECSA, then the **DSPNAME** field is labeled N/A.

Operands

CSMOWNID(*owner_identifier*)

Specifies the owner ID of the storage pool. The owner ID is the address space identifier (ASID). This is a required operand. The owner ID must be 2 - 4 hexadecimal characters in the form X'nn'. If less than 4 characters, the owner ID is padded on the left with zeros.

Use the following syntax as an alternative to the panel interface.

Syntax

```
▶▶—CSMOWNER— —CSMOWNID—(—owner_identifier—)—————▶▶
```

Equated symbol

Symbol

Description

IVTDSpace

The last CSM data space buffer processed

Sample output

```
CSMOWNER CSMOWNID(0002)
```

CSMOWNER Analysis

```
CSM 4K ECSA Pool 03B45000 for Owner ID 0002
EXT 03E20630 HDR 03E206B0 PHDR N/A BUF 03B35000 DSPNAME N/A
EXT 03E20630 HDR 03E20700 PHDR N/A BUF 03B36000 DSPNAME N/A
EXT 03E20630 HDR 03E20750 PHDR N/A BUF 03B37000 DSPNAME N/A
```



```

EXT 03E20630 HDR 03E207A0 PHDR N/A      BUF 03B38000 DSPNAME N/A
EXT 03E20630 HDR 03E207F0 PHDR N/A      BUF 03B39000 DSPNAME N/A
EXT 03E20630 HDR 03E20840 PHDR N/A      BUF 03B3A000 DSPNAME N/A
EXT 03E20630 HDR 03E20890 PHDR N/A      BUF 03B3B000 DSPNAME N/A
EXT 03E20630 HDR 03E208E0 PHDR N/A      BUF 03B3C000 DSPNAME N/A
EXT 03E20630 HDR 03E20930 PHDR N/A      BUF 03B3D000 DSPNAME N/A
EXT 03E20630 HDR 03E20980 PHDR N/A      BUF 03B3E000 DSPNAME N/A
EXT 03E20630 HDR 03E209D0 PHDR N/A      BUF 03B3F000 DSPNAME N/A
Number of primary buffers owned by 0002 in CSM 4K ECSA pool: 11
Number of _image_ buffers owned by 0002 in CSM 4K ECSA pool: 0

CSM 16K ECSA Pool 03B01000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 16K ECSA pool: 0
Number of _image_ buffers owned by 0002 in CSM 16K ECSA pool: 0

CSM 32K ECSA Pool 03AFE000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 32K ECSA pool: 0
Number of _image_ buffers owned by 0002 in CSM 32K ECSA pool: 0

CSM 60K ECSA Pool 03AFB000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 60K ECSA pool: 0
Number of _image_ buffers owned by 0002 in CSM 60K ECSA pool: 0

CSM 4K DSPACE31 Pool 032F3000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 4K DSPACE31 pool: 0
Number of _image_ buffers owned by 0002 in CSM 4K DSPACE31 pool: 0

CSM 16K DSPACE31 Pool 02AEB000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 16K DSPACE31 pool: 0
Number of _image_ buffers owned by 0002 in CSM 16K DSPACE31 pool: 0

CSM 32K DSPACE31 Pool 02AE5000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 32K DSPACE31 pool: 0
Number of _image_ buffers owned by 0002 in CSM 32K DSPACE31 pool: 0

CSM 60K DSPACE31 Pool 02ADF000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 60K DSPACE31 pool: 0
Number of _image_ buffers owned by 0002 in CSM 60K DSPACE31 pool: 0

CSM 180K DSPACE31 Pool 02ADC000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 180K DSPACE31 pool: 0
Number of _image_ buffers owned by 0002 in CSM 180K DSPACE31 pool: 0

CSM 4K DSPACE64 Pool 03AF8000 for Owner ID 0002
EXT 03B14500 HDR 03B14580 PHDR N/A      BUF 032F8000 DSPNAME CSM64001
EXT 03B14500 HDR 03B145D0 PHDR N/A      BUF 032F9000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14620 PHDR N/A      BUF 032FA000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14670 PHDR N/A      BUF 032FB000 DSPNAME CSM64001
EXT 03B14500 HDR 03B146C0 PHDR N/A      BUF 032FC000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14710 PHDR N/A      BUF 032FD000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14760 PHDR N/A      BUF 032FE000 DSPNAME CSM64001
EXT 03B14500 HDR 03B147B0 PHDR N/A      BUF 032FF000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14800 PHDR N/A      BUF 03300000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14850 PHDR N/A      BUF 03301000 DSPNAME CSM64001
EXT 03B14500 HDR 03B148A0 PHDR N/A      BUF 03302000 DSPNAME CSM64001
EXT 03B14500 HDR 03B148F0 PHDR N/A      BUF 03303000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14940 PHDR N/A      BUF 03304000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14990 PHDR N/A      BUF 03305000 DSPNAME CSM64001
EXT 03B14500 HDR 03B149E0 PHDR N/A      BUF 03306000 DSPNAME CSM64001
EXT 03B14500 HDR 03B14A30 PHDR N/A      BUF 03307000 DSPNAME CSM64001
Number of primary buffers owned by 0002 in CSM 4K DSPACE64 pool: 16
Number of _image_ buffers owned by 0002 in CSM 4K DSPACE64 pool: 0

CSM 16K DSPACE64 Pool 02AEE000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 16K DSPACE64 pool: 0
Number of _image_ buffers owned by 0002 in CSM 16K DSPACE64 pool: 0

```

CSMOWNER

```
CSM 32K DSPACE64 Pool 02AE8000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 32K DSPACE64 pool:      0
Number of _image_ buffers owned by 0002 in CSM 32K DSPACE64 pool:      0

CSM 60K DSPACE64 Pool 02AE2000 for Owner ID 0002
Number of primary buffers owned by 0002 in CSM 60K DSPACE64 pool:      0
Number of _image_ buffers owned by 0002 in CSM 60K DSPACE64 pool:      0
```

CSMPOOL

Use CSMPOOL to display the CSM control blocks for a specific size and type of CSM storage pool.

This command shows the pool control block, extents, and the list of ASIDs of the registered users of this pool. You can optionally request the primary headers of the registered users of the pool.

Operands

CSMTYPE(ECSA|DSPACE)

Specifies the type of storage to be displayed. The ECSA option displays storage from CSM extended common service area (ECSA). The DSPACE option displays storage from 31-bit backed and 64-bit backed CSM data space. This is a required operand.

CSMSIZE(4K|16K|32K|60K|180K)


Specifies the size of the storage pool to be displayed. This is a required operand.

DETAIL|NODETAIL

Specifies the level of detail to be displayed for the CSM storage pool. Enter DETAIL to display the entire contents of the CSM storage pool. If the DETAIL option is selected, the primary headers are also shown. The default is NODETAIL.

Use the following syntax as an alternative to the panel interface.

Syntax



Equated symbol

Symbol

Description

IVTDSPACE

The last CSM data space buffer processed

Sample output

```
CSMPOOL CSMSIZE(4K) CSMTYPE(DSPACE) DETAIL
```

```
CSMPOOL Analysis
```

```
CSM 4K DSPACE31 POOL
```

```

CSMPOOL: 032F3000
+0000 D7D6D6D3 20604000 00001000 00000001 | POOL.- .....
+0010 00000040 00000040 00000008 00000010 | ... ..
+0020 032F6500 02AF1A80 00000001 03ED0000 | .....
+0030 00000000 00000040 00000000 00000000 | .....
+0040 00000000 03E1B28C 03EC0000 00000100 | .....
+0050 00000040 06900001 032F6500 00000000 | ... ..
    
```

```

CSMEXT: 032F6500
+0000 C5E7E340 05800000 02AF2A80 00000000 | EXT .....
+0010 032F3000 01C42AF3 F0000800 02AF3000 | ....D.30.....
+0020 00000002 02AF3000 02B02FFF 00000010 | .....
+0030 00000010 00000000 00010000 00000010 | .....
+0040 F0C1C1C1 C3C3E2D4 00000000 00000000 | CSM31002.....
+0050 00000000 00000000 00000000 00000000 | .....
+0060 00000000 00000000 00000000 00000000 | .....
+0070 00000000 00000000 00000000 00000000 | .....
    
```

```

CSMHDR: 032F6580
+0000 C8C4D940 00500080 00000000 00000000 | HDR .&.....
+0010 00000000 02AF3000 032F6580 00000000 | .....
+0020 00000000 00000000 00000000 00000000 | .....
+0030 00000000 00000000 00000000 00000000 | .....
+0040 00000000 00000000 01000100 00000000 | .....
    
```

⋮

```

CSMEXT: 02AF2A80
+0000 C5E7E340 05800000 02AF2500 032F6500 | EXT .....
+0010 032F3000 01C42AF3 F0000800 02AF3000 | ....D.30.....
+0020 00000002 02B03000 02B12FFF 00000010 | .....
+0030 00000010 00000000 00010000 00000010 | .....
+0040 F0C1C1C1 C3C3E2D4 00000000 00000000 | CSM31002.....
+0050 00000000 00000000 00000000 00000000 | .....
+0060 00000000 00000000 00000000 00000000 | .....
+0070 00000000 00000000 00000000 00000000 | .....
    
```

```

CSMHDR: 02AF2B00
+0000 C8C4D940 00500080 00000000 00000000 | HDR .&.....
+0010 00000000 02B03000 02AF2B00 00000000 | .....
+0020 00000000 00000000 00000000 00000000 | .....
+0030 00000000 00000000 00000000 00000000 | .....
+0040 00000000 00000000 01000100 00000000 | .....
    
```

⋮

Users of CSM 4K DSPACE31 Pool:
0002

CSM 4K DSPACE64 POOL

```

CSMPOOL: 03AF8000
+0000 D7D6D6D3 20604020 00001000 00000001 | POOL.- .....
+0010 00000040 00000030 00000008 00000010 | ... ..
+0020 03B14500 032F6A80 00000001 03ED0020 | .....|.....
+0030 00000000 00000040 00000000 00000000 | .....
+0040 00000000 03E1B2A0 03EC2F60 00000100 | .....-.....
+0050 00000040 0B980001 03B14500 00000000 | ... .q.....
    
```

```

CSMEXT: 03B14500
+0000 C5E7E340 05800000 032F7A80 00000000 | EXT .....:.....
+0010 03AF8000 01C432F8 F0000800 032F8000 | ....D.80.....
    
```

CSMPOOL

```

+0020 00000001 032F8000 03307FFF 00000010 | ....."...... |
+0030 00000000 00000000 00010000 00000010 | ..... |
+0040 F0C1C1C1 C2C3E2D4 00000000 00000000 | CSM64001..... |
+0050 00000000 00000000 00000000 FFFF0000 | ..... |
+0060 00000000 00000000 00000000 00000000 | ..... |
+0070 00000000 00000000 00000000 00000000 | ..... |
:
CSMEXT: 032F7A80
+0000 C5E7E340 05800000 032F7500 03B14500 | EXT ..... |
+0010 03AF8000 01C432F8 F0000800 032F8000 | .....D.80..... |
+0020 00000001 03308000 03317FFF 00000010 | ....."...... |
+0030 00000010 00000000 00010000 00000010 | ..... |
+0040 F0C1C1C1 C2C3E2D4 00000000 00000000 | CSM64001..... |
+0050 00000000 00000000 00000000 00000000 | ..... |
+0060 00000000 00000000 00000000 00000000 | ..... |
+0070 00000000 00000000 00000000 00000000 | ..... |
:
CSMHDR: 032F7B00
+0000 C8C4D940 00500080 00000000 00000000 | HDR .&..... |
+0010 00000000 03308000 032F7B00 00000000 | .....#...... |
+0020 00000000 00000000 00000000 00000000 | ..... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 01000100 00000000 | ..... |
:
Users of CSM 4K DSPACE64 Pool:
0002

```

FINDDSIB

Use FINDDSIB to scan the ATCVT DSSIB queue for DSSIBs that meet specified selection criteria. The following information are displayed for each DSSIB selected:

- DSSIB address
- Procedure correlation identifier (PCID) of the request
- Owning SSCP name
- Real name of the destination logical unit
- Real network ID of the destination logical unit
- Alias name of the destination logical unit
- Alias network ID of the destination logical unit
- Adjacent SSCP in the originating direction

FINDDSIB has no required selection operands. If you enter no value for all selection operands, all DSSIBs are eligible for selection.

To select specific DSSIBs, you may enter a value for any of the selection operands below. All entered values must be present in the correct position within a DSSIB for it to be selected. For example, if you specify both a real name and an alias name, only DSSIBs with the specified real name in the RNAME position and the specified alias name in the ALIAS position are eligible for selection.

If you enter no value for a selection operand, DSSIBs with any value in that position are eligible for selection.

Use the Process positional operand to set the number of eligible DSSIBs that will actually be selected and displayed (all of them or just the first one encountered).

Operands

PCID

Specify 2–16 hexadecimal digits in the form X'x...'. for the PCID associated with the DSRLST request. Specify an even number of digits, otherwise the high-order 4 bits are assumed to be 0. If the PCID entered is fewer than 16 digits, then it is right-aligned, and a match occurs with all DSSIBs with PCIDs whose rightmost digits match the specified digits. The specified PCID is not padded with any characters.

Owning SSCP

The owning SSCP name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Real name

The real name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Real network ID

The real network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Alias name

The alias name of the DLU resource should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Alias network ID

The alias network ID of the DLU resource should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Adjacent SSCP

The adjacent SSCP in the originating direction should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Process


Use **First** to display the first DSSIB that meets the selection criteria. Otherwise, all DSSIBs that meet the selection criteria are displayed.

Routing in progress

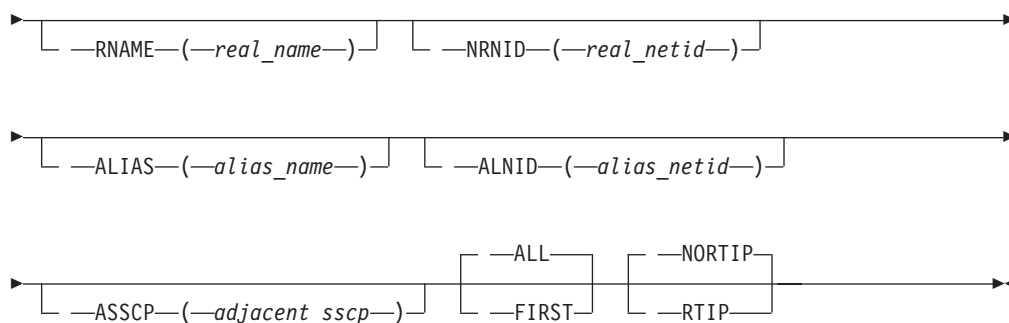
Use **Check RTIP** to display only DSSIBs that indicate “routing in progress” (“routing in progress” is indicated when bit DSSRTIP is on). Otherwise, FINDDSIB does not check for “routing in progress” (that is, the DSSRTIP bit is ignored).

Use the following syntax as an alternative to the panel interface.

Syntax



FINDDSIB



Sample output

FINDDSIB

FINDDSIB Analysis							
DSSIB	PCID	OSSCP	RNAME	NRNID	ALIAS	ALNID	ASSCP
05A13498	F0871BD0A7E3DF8C	XYZSCP05	TS0105	XYZNET	TS0105	XYZNET	ABCNET
DSSIBs processed:				1			
DSSIBs matching search criteria:				1			

FINDSIB

Use FINDSIB to scan a queue of SIBs for those that meet specified selection criteria. The following items are displayed for each SIB selected:

- SIB address
- Initiation finite state machine (SIBFSMIN)
- Termination finite state machine (SIBFSMTM)
- PLU NetID name
- PLU name
- SLU NetID name
- SLU name
- PLU network address
- SLU network address
- Procedure correlation identifier (PCID)

FINDSIB has no required operands. If you enter no values for all selection operands, all SIBs on the ATCVT SIB queue are eligible for selection.

To scan an SIB queue other than the ATCVT SIB queue (such as the primary or secondary SIB queue off of an RDTE), you must specify a primary or secondary SIB queue. Specify only one queue. If both a primary queue and a secondary queue are specified, only the secondary queue will be used.

To select specific SIBs, enter a value for any of the selection operands below. All values entered must be present in the correct position within the SIB for it to be selected. If you select both PLU Name and PCID, only SIBs with the specified PLU Name in the PLUNAME position and the specified PCID in the PCID position are eligible for selection.

Note: You might need to find SIBs for a resource but do not know whether the resource is the PLU or SLU. In this special case, you can specify the resource name

for both the PLU name and the SLU name, and if the resource name is found in either one, a match occurs. The SES function can also be used to find all sessions for a specified resource name.

If you enter no value for a selection operand, SIBs with any value in that position are eligible for selection from the specified SIB queue.

Use the Process operand to set the number of eligible SIBs that will actually be selected and displayed (all of them or just the first one encountered).

Operands

PLU name

The PLU name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

PLU NetID

The PLU network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

SLU name

The SLU name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

SLU NetID

The SLU network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

PCID

Specify 2–16 hexadecimal digits in the form X'x...' for the PCID. Specify an even number of digits, otherwise the high-order 4 bits are assumed to be 0. If the PCID entered is fewer than 16 digits, then it is right-aligned, and a match occurs with all SIBs with PCIDs whose rightmost digits match the specified digits. The specified PCID is not padded with any characters.

PLU network address

Specify 1–12 hexadecimal digits in the form X'x...' for the network address of the PLU. If you specify fewer than 12 digits, FINDSIB selects a network address whose rightmost digits match the specified digits.

Example: Subarea 12 Input = 1204BC Dump Data = 0000001204BC
 Element 04BC

SLU network address

Specify 1–12 hexadecimal digits in the form X'x...' for the network address of the SLU. If you specify fewer than 12 digits, FINDSIB selects a network address whose rightmost digits match the specified digits.

Example: Subarea A Input = A1123 Dump Data = 000000A0123
 Element 0123

Note: The following three operands, Displacement 1, Value 1, and Value 1 Type, must be specified together. They allow any field in an SIB to be checked for a user-specified value.

Displacement 1

Enter the displacement into the SIB where Value 1 is to be found. The maximum decimal displacement is 4095, and the maximum hexadecimal displacement is X'FFF'.

Value 1

Only SIBs containing this character, hex, or binary value at the displacement specified in Displacement 1 are selected.

Value may contain character or hexadecimal data of 1–8 bytes in length. Hexadecimal data should contain an even number of up to 16 hexadecimal digits in the form X'xx...', otherwise the high order 4 bits are assumed to be 0.

Binary data can be used to look at a particular bit within a byte. You may specify 1 byte of binary data in the form X'xx'. Only 1 bit within the byte may be selected. Therefore, you can specify only the following hexadecimal values: 01, 02, 04, 08, 10, 20, 40, and 80. A value with more than 1 bit set (for example, 82) will not be accepted. If you want to test 2 bits within the same byte, you must use Displacement 2, Value 2, and Value 2 Type, as well as Displacement 1, Value 1, and Value 1 Type.

Value 1 Type

Enter B for binary, C for character, or X for hexadecimal to indicate the type of data entered for Value 1.

Note: The following three operands, Displacement 2, Value 2, and Value 2 Type, are used together.

Displacement 2

Same as **Displacement 1**.

Value 2

Same as **Value 1**.

Value 2 Type

Same as **Value 1 Type**.

Note: If both (Displacement 1, Value 1, Value 1 Type) and (Displacement 2, Value 2, Value 2 Type) are specified, both sets of conditions must be met for a SIB to be selected.

You may specify only one queue, Primary SIB, or Secondary SIB.

Primary SIB

Enter the address of an SIB on the primary SIB queue off of an RDTE. The address must be 1–8 hexadecimal digits in the form X'x...'. If the address specified is fewer than eight digits, it is padded on the left with zeros.

Secondary SIB

Enter the address of an SIB on the secondary SIB queue off of an RDTE. The address must be 1–8 hexadecimal digits in the form X'x...'. If the address specified is fewer than eight digits, it is padded on the left with zeros.

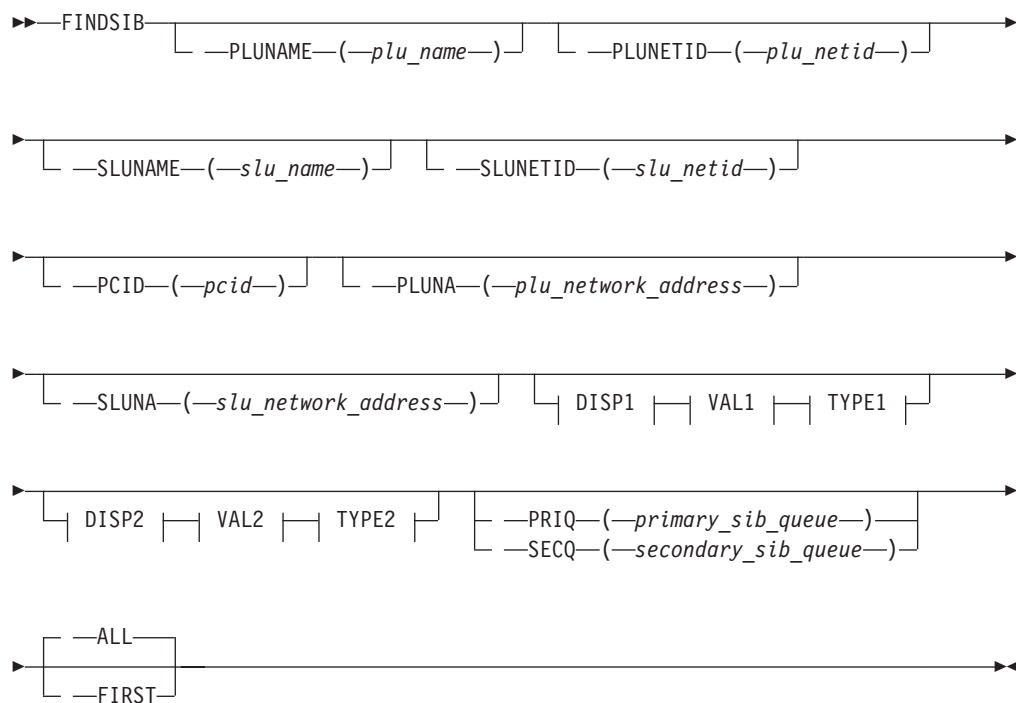
Note: If All is specified for Primary SIB or Secondary SIB, all elements from the first match are considered regardless of the address.

Process

Use **First** to display the first SIB that meets the selection criteria. Otherwise, all SIBs that meet the selection criteria are displayed.

Note: Scanning the entire SIB queue can take a long time.

Use the following syntax as an alternative to the panel interface.

Syntax**DISP1**

```
| -DISP1 ( -displacement ) |
```

VAL1

```
| -VAL1 ( -data_value ) |
```

TYPE1

```
| -TYPE1 ( -data_type ) |
```

DISP2

```
| -DISP2 ( -displacement ) |
```

VAL2

```
| -VAL2 ( -data_value ) |
```

TYPE2

```
| -TYPE2 ( -data_type ) |
```

FINDSIB

Sample output

FINDSIB PLUNETID(NETB) PLUNAME(ECHOB1B) SLUNETID(NETC) SLUNAME(C01D0067)

FINDSIB Analysis

SIB	ADDR	FSMS	PLUNETID	PLUNAME	SLUNETID	SLUNAME	PLUNA	SLUNA	PCID
069AB830	3C00	NETB	ECHOB1B	NETC	C01D0067	000000000000	000000000000	ECC39EEE2A54E5D9	

SIBs processed: 1095
SIBs matching search criteria: 1

FNDADJCP

FNDADJCP scans all of the partner nodes that have CP-CP sessions with this host for the given resource. If a resource is not provided, all partner nodes are displayed.

FNDADJCP has no required operands. If you do not enter a resource name, all ACPCB control blocks are formatted.

Operands

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Resource name

The resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Use the following syntax as an alternative to the panel interface.

Syntax



NETID

`—NETID—(—netid—)`

RESNAME

`—RESNAME—(—resource_name—)`

Sample output

FNDADJCP

FNDADJCP Analysis

ACPCB:	069EF200					ACPNETZ	A04P
+0000	40C1C3D7	D5C5E3E9	40404040	C1F0F4D7	883A.....4}		
+0010	F8F8F3C1	00040008	00000000	069EF4D0		
+0020	00000000	00000000	00000000	00000000		
+0030	00000000	00000000	00000000	00000000		
+0040	00000000	00000000	00000000	4000000C		
+0050	00000000	00000000	4000000C	00000000		

```

+0060 00000000 00000000 10000000 | ..... |
:
ACPCB: 069EF4D0
+0000 40C1C3D7 D5C5E3C1 40404040 C1F0F4D7 | ACPNETA A04P |
+0010 F8F8F7C1 00040008 069EF200 069EF458 | 887A.....2...4. |
+0020 00000000 00000000 00000000 00000000 | ..... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 00000000 4000000C | ..... |
+0050 00000000 00000000 4000000C 00000000 | ..... |
+0060 00000000 00000000 10000000 | ..... |

```

FNDANDCB

Use FNDANDCB to help diagnose problems with CP-CP sessions between this host and adjacent nodes. For a particular resource, FNDANDCB finds and formats the IstandCB and ISTCPCAP control blocks.

Operands

Resource name

The resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.


The resource name is required.

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Use the following syntax as an alternative to the panel interface.

Syntax



RESNAME



NETID



Sample output

FNDANDCB RESNAME(A04P887A)

FNDANDCB Analysis

```

ANDCB: 069C6B10
+0000 C1D5C3C2 D5C5E3C1 4BC1F0F4 D7F8F8F7 | ANCBNETA.A04P887 |
+0010 C1404040 40000000 00000000 C3D7E2E5 | A .....CPSV |
+0020 C3D4C740 00000000 00000000 58588000 | CMG ..... |
+0030 069C6A24 00000000 00000000 00C1BE88 | ..].....A_h |
+0040 04000000 00000000 00000000 00000000 | ..... |
+0050 00000000 00000000 00000000 00000000 | ..... |
+0060 00000000 00000000 00000000 00000000 | ..... |

```

FNDANDCB

```

+0070 00000000 00000000 00000000 A6C1C094 | .....wA{m
+0080 33E74905 02000000 00000000 069C69F8 | .X.....8
+0090 00000000 00000000 00000000 00000000 | .....
+00A0 06B72EE0 89400000 02000000 00000000 | ... \i .....
+00B0 00000000 00000000 00000000 80000000 | .....
+00C0 00000001 10F01002 02000000 06B6DD00 | .....0.....
+00D0 01000000 00000000 00000000 00000000 | .....
+00E0 00000000 80000000 00000000 A0801010 | .....
+00F0 00000000 00000000 0679ABE8 00000000 | ..... ^ .Y....
+0100 00000000 00000000 00000000 00          | .....

CPCAP: 0679ABE8
+0000 C3D7C3C1 000C12C1 00000000 80800000 | CPCA...A..... |

```

FNDCOS

Use FNDCOS to format mode tables, mode table entries, and Class of Service entries found in those mode table entries.

FNDCOS formats and displays the following control blocks:

- ISTCSTRU
- ISTMCOSS
- ISTMDTAB
- ISTNDWED
- ISTTGWGT

In order to reduce repetitious output, the control blocks ISTCSTRU, ISTNDWED, and ISTTGWGT will not display for consecutive, identical Class of Service names.

FNDCOS has no required operands.

Operands

Mode table

The name of the mode table should be 1–8 alphanumeric characters.

Mode name


The name of the entry in the mode table should be 1–8 alphanumeric characters.

Class of Service name

The name of the entry in the APPN Class of Service table should be 1–8 alphanumeric characters.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

FNDCOS

```

                                COS Analysis

MDTAB: 0695A460
+0000 C5D4E2D4 D6C4C540 00000015 C0000014 | EMSMODE ....{... |
+0010 0695A9A0 0695A4A0 0695A4A0 00000800 | .nz..nu..nu..... |
+0020 00000000 0695A0E0 | .....n.\ |

MCOSS: 0695A9A0
+0000 C4E8D5C1 D4C9C340 7BC3D6D5 D5C5C3E3 | DYNAMIC #CONNECT |
+0010 0693A720 00000000 0695A960 00000000 | .lx.....nz-.... |
+0020 00000000 00000000 | ..... |

CSTRU: 0693A720
+0000 7BC3D6D5 D5C5C3E3 0693A2E0 0693AB60 | #CONNECT.lxs\l.- |
+0010 00000008 40000000 0693A760 0693A920 | .... .lx-.lz. |
+0020 40000000 0693A960 0693AB20 06A4E080 | ....lz-.l...u\ |
+0030 40 |

TGWGT: 0693A760
+0000 00000000 0693A7A0 00000000 01C0004C | .....lx.....{.< |
+0010 75FF00FF 00FF00FF 1E | ..... |
:
.
.
TGWGT: 0693A920
+0000 0693A8E0 00000000 00FF00FF 01C000FF | .ly\.....{.. |
+0010 00FF00FF 00FF00FF F0 | .....0 |

NDWED: 0693A960
+0000 00000000 0693A9A0 00001F05 | .....lz..... |
:
.
.
NDWED: 0693AB20
+0000 0693AAE0 00000000 4000FFA0 | .l.\.... ... |

```

FNDDECB

Use FNDDECB to format a directory entry and its parent directory entries.

FNDDECB formats and displays the ISTDECB control block.

Operands

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Note: If you do not specify a network ID, the host network ID will be used to form a fully qualified network name.

Resource name

The resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

The resource name is required.

Use the following syntax as an alternative to the panel interface.

FNDDECB

Syntax

►► FNDDECB ————— —RESNAME—(—resource_name—)————►
└── —NETID—(—netid—) ─┘

Sample output

FNDDECB RESNAME(A44N)

DECB Analysis

```
DECB: 068E9158
+0000 C4C5C3C2 44000004 06BCB7D8 068E9498 | DECB.....Q..mq |
+0010 00010000 C1F4F4D5 40404040 D5C5E3C1 | ...A44N NETA |
+0020 40404040 000400F4 07041548 067D7968 | ...4.....'` |
+0030 06A4CFC0 A682A447 068E90F0 40000048 | .u.{wbu....0 ... |
+0040 06A38BE8 06A63978 06A38B80 06A62D88 | .t.Y.w...t...w.h |
:
```

Parent DECB chain

```
DECB: 068E90F0
+0000 C4C5C3C2 44000004 06A64E58 06A63B80 | DECB....w+..w.. |
+0010 00000000 C1F0F2D5 40404040 D5C5E3C1 | ...A02N NETA |
+0020 40404040 000400F6 00000000 06CEE788 | ...6.....Xh |
+0030 06A49890 A682A461 00000000 40000048 | .uq.wbu/.... ... |
+0040 06A38B80 06EE8B80 00000000 00000000 | .t..... |
+0050 06C908A8 06C90AB0 00000000 | .I.y.I..... |
```

FNDENDEL

Use FNDENDEL to help diagnose problems with adjacent end nodes. For a particular resource, FNDENDEL provides the associated ENDEL control block.

FNDENDEL has no required operands. If you do not enter a resource name, all ENDEL control blocks are formatted.

Operands

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Resource name

The resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Use the following syntax as an alternative to the panel interface.

Syntax

►► FNDENDEL —————
└── NETID ─┘ └── RESNAME ─┘

NETID

└── —NETID—(—netid—) ─┘

RESNAME

┌─── RESNAME──(—resource_name—)──┐

Sample output

FNDENDEL RESNAME(A04P208A)

FNDENDEL Analysis

```
ENDEL: 06A41020
+0000 C5D5C4D3 00000000 06A41048 000DD5C5 | ENDL.....u....NE |
+0010 E3C14BC1 F0F4D7F2 F0F8C140 40404000 | TA.A04P208A . |
```

FNDLCB

Use FNDLCB to help diagnose problems with directory search requests. For a particular procedure correlation identifier (PCID), FNDLCB finds and formats the following control blocks:

- LCB
- LCB extension
- OSCB
- Original and best reply PLOCBs from the queue of LCB control blocks

FNDLCB has no required operands. If you do not enter a PCID, all control blocks from the previous list are formatted.


Operands

PCID

Specify 1–16 hexadecimal digits in the form X'x...'. for the PCID. Specify an even number of digits, otherwise the high-order 4 bits are assumed to be 0. If the PCID entered is fewer than 16 digits, then it is right-aligned, and a match occurs with all LCBs with PCIDs whose rightmost digits match the specified digits. The specified PCID is not padded with any characters.

Use the following syntax as an alternative to the panel interface.

Syntax



PCID

┌─── PCID──(—pcid—)──┐

Sample output

FNDLCB

FNDLCB Analysis

```
LCB: 06BE4E90
+0000 40D3C3C2 00000000 06BE4170 00000000 | LCB....._..... |
+0010 00000000 00000000 00000000 00000000 | ..... |
```

FNDLCB

```
+0020 D3A3D286 D58FA88F 000DD5C5 E3C34BC3 | LtKfN.y...NETC.C |
+0030 F0F4D7F2 F0F8C100 00000000 00000000 | 04P208A..... |
+0040 00000002 11100000 00000000 00000000 | ..... |
+0050 21110000 00000002 11000000 00000000 | ..... |
:
:
Best reply
PLOCB: 06CE3B00
+0000 40D7D3D6 00000000 00000000 00000000 | PLO..... |
+0010 00000000 00000000 00000000 00000000 | ..... |
+0020 00000000 00808000 00000000 00000000 | ..... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 A4000040 00000000 D3A3D286 | ....u.. ....LtKf |
:
:
Original reply
PLOCB: 06BBC2D8
+0000 40D7D3D6 00000000 00000000 00000000 | PLO..... |
+0010 00000000 00000000 00000000 00000000 | ..... |
+0020 00000000 40C4A000 000008B6 00000000 | .... D..... |
+0030 0004C1F0 F1D50000 00000004 D5C5E3C1 | ..A01N.....NETA |
+0040 00000000 A0000080 00002110 D3A3D286 | .....LtKf |
+0050 D58FA88F 000DD5C5 E3C34BC3 F0F4D7F2 | N.y...NETC.C04P2 |
:
:
LCBEXT: 06C6E178
+0000 D3C3C2C5 10000000 0004D5C5 E3C10000 | LCBE.....NETA.. |
+0010 00000000 0008C1F3 F1C9F4F8 F9F20000 | .....A31I4892.. |
+0020 0004D5C5 E3C10000 00000000 0008C1F3 | ..NETA.....A3 |
+0030 F1D7F4F8 F9C10000 0004D5C5 E3C10000 | 1P489A.....NETA.. |
+0040 00000000 0004C1F0 F2D50000 00000000 | .....A02N..... |
+0050 00000000 00C00000 068F8A38 00000000 | .....{..... |
:
:
LCB: 06BE4170
+0000 40D3C3C2 06BE4E90 06BE4560 00000000 | LCB._+...-.... |
+0010 00000000 00000000 00000000 00000000 | ..... |
+0020 CFA9CD86 D295A38D 000DD5C5 E3C34BC3 | .z.fKnt_...NETC.C |
+0030 F0F4D7F1 F6F5C100 00000000 00000000 | 04P165A..... |
+0040 00000002 01100000 00000000 00000000 | ..... |
+0050 20110000 00000002 01000000 00000000 | ..... |
:
:
```

FNDNDREC

Use FNDNDREC to help diagnose topology and routing problems.

FNDNDREC scans the topology and route selection database for node records matching the given resource for the SINGLE NODE option and formats the NDREC control block.

FNDNDREC also provides summary information output of user-selected criteria.

Operands

Resource name

The resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

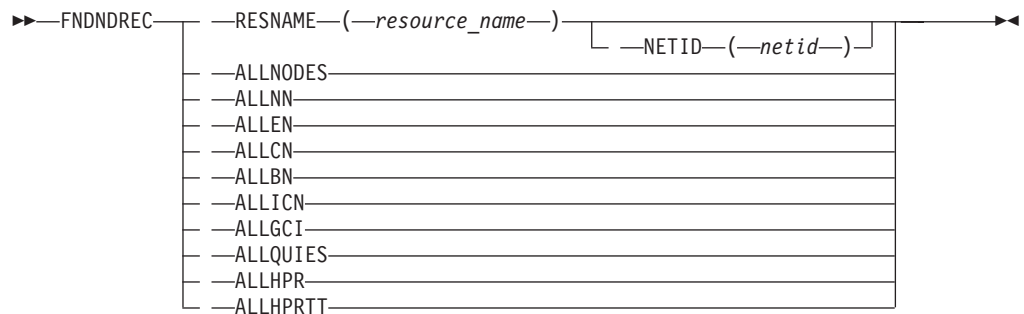
ALLNODES

Displays a summary of all node records.

- ALLNN**
Displays a summary of all network nodes.
- ALLEN**
Displays a summary of all end nodes.
- ALLCN**
Displays a summary of all connection network nodes.
- ALLBN**
Displays a summary of all border nodes.
- ALLICN**
Displays a summary of all interchange nodes.
- ALLGCI**
Displays a summary of all nodes with GCI on.
- ALLQUIES**
Displays a summary of all nodes with quiescing on.
- ALLHPR**
Displays a summary of all nodes with base HPR on.
- ALLHPRTT**
Displays a summary of all nodes with HPRTT on.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

FNDNDREC RESNAME(A04P883A) NETID(NETZ)

```

                                FNDNDREC Analysis

NDREC: 068F3A48
+0000 D5C4D9C3 0000002B 000DD5C5 E3E94BC1 | NDRC.....NETZ.A
+0010 F0F4D7F8 F8F3C140 40404000 40000004 | 04P883A . ...
+0020 00000000 00000000 40000004 00000000 | .....
+0030 00000000 00000000 00000000 00000000 | .....
+0040 00000000 00000000 00000000 00000000 | .....
+0050 00000F00 00000001 00000000 00000000 | .....
+0060 00000000 00000000 00000000 00000000 | .....
+0070 00000000 00000000 12440DD5 C5E3E94B | .....NETZ.
+0080 C1F0F4D7 F8F8F3C1 00000000 00000000 | A04P883A.....
+0090 00000000 00000000 00000000 00000C45 | .....
+00A0 0A800000 0000FF60 13000000 00000000 | .....
+00B0 00000000 00000000 00000000 0000      | .....
  
```

FNDNDREC

FNDNDREC ALLNODES

FNDNDREC Analysis

Node Records Summary

Node Name	NDRECAAddr	NodeType	Time	SeqNo	GCI	QUIES	HPR
NET000A.M001G	14B6F2B0	EN	5	00000000	N	N	NO
NET000A.E502CDRM	14B6F010	EN	15	0000000C	N	N	CONTR
NET000A.CIESB049	14B6F470	NN	5	00000000	N	N	NO
NET000A.CIESB050	14B6F550	NN	5	00000000	N	N	NO
NET000A.M802P	14B6F1D0	NN	5	00000000	N	N	NO
NET000A.M802O	14B6F390	NN	5	00000000	N	N	NO
Total Number of Nodes found:		6					

FNDNDWGT

Use FNDNDWGT to help diagnose topology and routing problems.

FNDNDWGT scans the topology and route selection database for node records matching the given resource name. It will go through the node weight control blocks to determine the node weight.

Operands

Origin resource name

The origin resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

The origin resource name is required.

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Class of Service name

The name of the entry in the APPN Class of Service table should be 1–8 alphanumeric characters. If cosname is not specified, the default cosname is #CONNECT.

Use the following syntax as an alternative to the panel interface.

Syntax

```
►► FNDNDWGT —ORESNAME—(—origin_resource_name—) [ —NETID—(—netid—) ] [ —COSNAME—(—cosname—) ] ◀◀
```

Sample output

```
FNDNDWGT NETID(NET000A) ORESNAME(E502CDRM) COSNAME(#CONNECT)
```

FNDNDWGT Analysis

NodeName	NDRECAAddr	WEIGHT
NET000A.E502CDRM	14B6F010	60

FNDNODE

Use FNDNODE to format one or more APPN adjacent end nodes or adjacent network nodes, or both.

FNDNODE formats and displays the ISTAENCB and ISTANNCB control blocks.

FNDNODE has no required operands.

Operands

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Note: If you specify a resource name but do not specify a network ID, the host network ID will be used to form a fully qualified network name.

Resource name

The resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Type

Enter BOTH to format both end nodes and network nodes. BOTH is the default. Enter EN to format only end nodes. Enter NN to format only network nodes.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

FNDNODE

```

                                FNDNODE Analysis
                                End Node List

AENCB: 06B24708
+0000 C5D5C3C2 00000000 06B246A0 00000000 | ENCB.....
+0010 00000000 00000000 000DD5C5 E3C24BC2 | .....NETB.B
+0020 F0F1D7F0 F0F9C140 40404000 01000074 | 01P009A .....
+0030 0100006E 40000000 00000000 00000000 | ...> .....
+0040 00000000 00000000 00000000 00000000 | .....
+0050 06B094D0 00000000 00000000 00000000 | ..m}.....
:
AENCB: 06A4A088
+0000 C5D5C3C2 06A4A0F0 00000000 06A4A5D0 | ENCB.u.0.....uv}
+0010 06A4A6A0 00010000 000DD5C5 E3E84BC1 | .uw.....NETY.A
+0020 F0F3D7F8 F8F2C140 40404000 00000000 | 03P882A .....
+0030 00000000 40000000 00000000 00000000 | ....
+0040 00000000 00000000 00000000 00000000 | .....
+0050 068F34D0 00000000 00000000 00000000 | ...}.....
  
```


Parent: NETA.A81N COSNAME: #CONNECT

```
REREC: 06A4F5F8
+0000 E3D9C5C3 00000000 068F30E8 06A4D0B0 | TREC.....Y.u}. |
+0010 00000000 06A4F440 00000000 00000000 | .....u4 ..... |
+0020 00000000 00010000 06A4F440 00000000 | .....u4 .... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 00000000 00000000 | ..... |
```

Sibling chain

NAME: NETA.A500N

```
REREC: 06A4F440
+0000 E3D9C5C3 06A4F5F8 068F3020 06A4D050 | TREC.u58....u}& |
+0010 00000000 00000000 06A3C480 00D20005 | .....tD..K.. |
+0020 000000D7 00020000 00000000 00000000 | ...P..... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 00000000 00000000 | ..... |
```

No children

FNDSCCB

Use FNDSCCB to format all ISTLCBs for a specific search concentration control block.

FNDSCCB formats and displays the ISTLCB control block.

Operands

Network ID

The network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Note: If you do not specify a network ID, the host network ID will be used to form a fully qualified network name.


Resource name

The resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

The resource name is required.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

FNDSCCB RESNAME(B01N)

FNDSCCB Analysis

```
LCB: 075E2D40
+0000 40D3C3C2 075BB6B0 075E22C0 07258020 | LCB.$...;.{.... |
+0010 075E22C0 00000000 0751C640 07521000 | .;.{.....F .... |
+0020 E7F3A765 0E1691B0 0009D5C5 E3C24BC2 | X3x...j...NETB.B |
```

FNDSCCB

+0030	F0F1D500	00000000	00000000	00000000	01N.....
+0040	00020001	20000000	00000000	00000000
+0050	00020002	00020000	20000000	00000000
+0060	00000000	00020000	00020001	20000000
+0070	00000000	00000000	01000067	0004C2F0BO
+0080	F1D50000	00000004	D5C5E3C2	00000000	1N.....NETB....
+0090	80000000	D6E2C3C2	00000000	00000000	...OSCB.....
+00A0	00000000	00000000	075E2D40	1C000000;.
+00B0	0005C1F5	F0F0D500	00000004	D5C5E3C1	..A500N.....NETA
+00C0	00000000	00000000	00000000	00000000
+00D0	13280000	00000000	071BAB00	00000000
+00E0	00000000	00000000	00000000	00000000
+00F0	00000000	00000000	00000000	00000000

FNDSITCB

Use FNDSITCB to help diagnose problems with the session services for LU-LU sessions. For a particular procedure correlation identifier (PCID), PLU name or network identifier, or SLU name or network identifier, FNDSITCB provides the associated SITCB control block.

FNDSITCB has no required operands. If you enter no values for all selection operands, all SITCB control blocks on the queue are eligible for selection.

Operands

PLU name

The PLU name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

PLU network ID

The name representing the network ID of another network outside the host network where a resource resides should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

SLU name

The SLU name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

SLU network ID

The name representing the network ID of another network outside the host network where a resource resides. SLU Network ID should be 1–8 alphanumeric characters. If it contains fewer than eight characters, the leftmost characters are compared.


PCID

Specify 1–16 hexadecimal digits in the form X'x...'. Specify an even number of digits, otherwise the high-order 4 bits are assumed to be 0. If the PCID entered is fewer than 16 digits, then it is right-aligned, and a match occurs with all SITCBs with PCIDs whose rightmost digits match the specified digits. The specified PCID is not padded with any characters.

Note: You might need to find SITCBs for a resource but do not know whether the resource is the PLU or SLU. In this case, you can specify the resource name for both the PLU name and the SLU name, and if the resource name is found in either one, a match occurs.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

FNDSITCB

FNDSITCB Analysis

SITCB: 06ACBAA0					
+0000	E2C9C3C2	00000000	00D7D5E2	04000000	SICB.....PNS....
+0010	00000000	00000000	86881220	00000000fh.....
+0020	00000000	00000000	06ACB020	FCD9C3C6RCF
+0030	FCE2D9D8	FCE2D9D8	00000000	FCE2D9D8	.SRQ.SRQ.....SRQ
+0040	00000000	00000000	00000000	00000000
+0050	00000000	00000000	00000000	00000000
:					
SITCB: 06ACB020					
+0000	E2C9C3C2	00000000	00D7D5E2	04000000	SICB.....PNS....
+0010	00000000	00000000	86881220	00000000fh.....
+0020	00000000	06ACBAA0	06C24AA0	FCD9C3C6B+..RCF
+0030	FCE2D9D8	FCE2D9D8	00000000	FCE2D9D8	.SRQ.SRQ.....SRQ
+0040	00000000	00000000	00000000	00000000
+0050	00000000	00000000	00000000	00000000
:					

FNDTGREC

Use FNDTGREC to help diagnose topology and routing problems. For an origin control point (CP), FNDTGREC with the DETAIL option formats the NDREC and TGREC control blocks linked between it and the destination control point.

The name of the destination CP is in the TGREC control block. The name of the origin CP is in the NDREC control block and is a required operand for FNDTGREC. The CP name is in the form of a network identifier and a resource name.

FNDTGREC, with the different summary options, provides the formatted origin NDREC control block and the TG records summary.

Operands

Origin resource name

The origin resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

The origin resource name is required.

Origin network ID

The origin network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Destination network ID

The destination network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Destination resource name

The destination resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

ALLOPER

Displays a summary of all TGs with OPER status.

ALLINOP

Displays a summary of all TGs with INOP status.

ALLGCITG

Displays a summary of all TGs with GCI on.

ALLQUITG

Displays a summary of all TGs with quiescing on.

ALLICTG

Displays a summary of all interchange TGs.

ALLENTG

Displays a summary of all endpoint TGs.

ALLBNTG

Displays a summary of all TGs with the border node indicator on.

ALLICLTG

Displays a summary of all ICL TGs.

ALLIRTG

Displays a summary of all intermediate routing TGs.

ALLBEXTG

Displays a summary of all branch extender TGs.

ALLHPTG

Displays a summary of all TGs with base HPR only.

ALLHPTTG

Displays a summary of all TGs with HPR Tower.

ALLODDTG

Displays a summary of all TGs with ODD sequence number.

ALLUSATG

Displays a summary of all usable TGs.

ALLTG

Displays a summary of all TGs.

Use the following syntax as an alternative to the panel interface.

Syntax



ORESNAME

|— —ORESNAME—(—origin_resource_name—)|

ONETID

|— —ONETID—(—origin_netid—)|

DNETID

|— —DNETID—(—destination_netid—)|

DRESNAME

|— —DRESNAME—(—destination_resource_name—)|

TG Summary

—ALLOPER—
—ALLINOP—
—ALLGCITG—
—ALLQUITG—
—ALLICTG—
—ALLENTG—
—ALLBNTG—
—ALLICLTG—
—ALLIRTG—
—ALLBEXTG—
—ALLHPRTG—
—ALLHPTTG—
—ALLODDTG—
—ALLUSATG—
—ALLTG—

Sample output

FNDTGREC ORESNAME(N317408) ONETID(NETA)

FNDTGREC Analysis

NDREC: 068F37F0					
+0000	D5C4D9C3	00000024	000CD5C5	E3C14BD5	NDRC.....NETA.N
+0010	F3F1F7F4	F0F84040	40404000	40000004	317408
+0020	06A08560	06A08800	40000004	00000000	..e-..h.
+0030	00000000	00000000	00000000	00000000
+0040	00000000	00000000	00000000	00000000
+0050	00000F40	00000001	00000000	00000000
+0060	00000000	00000000	00000000	00000000
+0070	00000000	00000000	11440CD5	C5E3C14BNETA.
+0080	D5F3F1F7	F4F0F800	00000000	00000000	N317408.....
+0090	00000000	00000000	00000000	00000C45
+00A0	0A800000	00468000	23000000	00000000
+00B0	00000000	00000000	00000000	0000
TGREC: 06A08560					
+0000	E3C7D9C3	00000000	06A08640	0F700000	TGRC.....f
+0010	40000004	00000000	00000000	00000024
+0020	068F37F0	068F3020	069E6D40	00000001	...0....._

FNDTGREC

```

+0030 00000000 00000000 00000000 00000000 .....
+0040 00000000 00000000 00000000 00000000 .....
+0050 00000000 00000000 16470000 00348076 .....
+0060 00000000 00000000 204C0000 00000000 .....<.....
+0070 00000000 00000000 00000000 00000000 .....
+0080 00000000 00001446 12800109 D5C5E3C1 .....NETA
+0090 4BC1F0F1 D5000000 00040000 00000000 .A01N.....
+00A0 00000000 00000000 00000000 00000000 .....
+00B0 00000000 00000000 00000000 00000000 .....
+00C0 00000000 00000000 00000000 00000000 .....
+00D0 00000000 00000000 .....

```

TGREC: 06A08640

```

+0000 E3C7D9C3 06A08560 06A08720 0F700000 TGRC..e-.g....
+0010 40000004 00000000 00000000 00000024 .....
+0020 068F37F0 068F30E8 06A07020 00000001 ...0...Y.....
+0030 00000000 00000000 00000000 00000000 .....
+0040 00000000 00000000 00000000 00000000 .....
+0050 00000000 00000000 16470000 00260076 .....
+0060 00000000 00000000 204C0000 00000000 .....<.....
+0070 00000000 00000000 00000000 00000000 .....
+0080 00000000 00001446 12800109 D5C5E3C1 .....NETA
+0090 4BC1F0F2 D5000000 01360000 00000000 .A02N.....
+00A0 00000000 00000000 00000000 00000000 .....
+00B0 00000000 00000000 00000000 00000000 .....
+00C0 00000000 00000000 00000000 00000000 .....
+00D0 00000000 00000000 .....

```

⋮

FNDTGREC ORESNAME(E502CDRM) ALLOPER

FNDTGREC Analysis

NDREC: 14B6F010

```

+0000 D5C4D9C3 00000149 0010D5C5 E3F0F0F0 NDRC.....N
+0010 C14BC5F5 F0F2C3C4 D9D44000 40000004 A.E502CDRM
+0020 14CF8010 14CF8E10 40000004 14B76038 .....
+0030 14B76038 00000000 00000000 00000000 ..-.....
+0040 00000000 00000000 00000000 00000000 .....
+0050 00000F08 00000001 00000000 00000000 .....
+0060 00000000 00000000 00000000 00000000 .....
+0070 00000005 00000000 40000004 14B6F0F0 .....
+0080 14B6F0F0 14B6F0FE 00000000 00000000 ..00..0....
+0090 154410D5 C5E3F0F0 F0C14BC5 F5F0F2C3 ...NET000A.
+00A0 C4D9D400 00000000 00000000 00000000 DRM.....
+00B0 00000000 00000C45 0A800000 000C8068 .....
+00C0 13180000 00000000 00000000 00000000 .....
+00D0 00000000 0000 .....

```

TG Records Summary

Dest CPNAME	TGRECAdr	TGN	STAT	Time	GCI	QUI	HPR/T	BEX	ICL	BN	TGTYPE
NET000A.M802P	14CF8010	22	OPER	15	N	N	N/N	N	N	N	ENDPT
NET000A.M802P	14CF81D0	21	OPER	15	N	N	N/N	N	N	N	ENDPT
NET000A.M8020	14CF8550	21	OPER	15	N	N	N/N	N	N	N	ENDPT
NET000A.M8020	14CF8710	22	OPER	15	N	N	N/N	N	N	N	ENDPT
NET000A.CIESB049	14CF88D0	22	OPER	15	N	N	Y/Y	N	N	N	ENDPT
NET000A.CIESB050	14CF8A90	21	OPER	15	N	N	Y/Y	N	N	N	ENDPT
NET000A.CIESB050	14CF8C50	22	OPER	15	N	N	Y/Y	N	N	N	ENDPT
NET000A.CIESB049	14CF8E10	21	OPER	15	N	N	Y/Y	N	N	N	ENDPT

Total Number of TG Records found: 8

FNDTGWGT

Use FNDTGWGT to help diagnose topology and routing problems.

Operands

Origin resource name

The origin resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

The origin resource name is required.

Origin network ID

The origin network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Destination network ID

The destination network ID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Destination resource name

The destination resource name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, the leftmost characters are compared.

Class of Service name

The name of the entry in the APPN Class of Service table should be 1–8 alphanumeric characters.

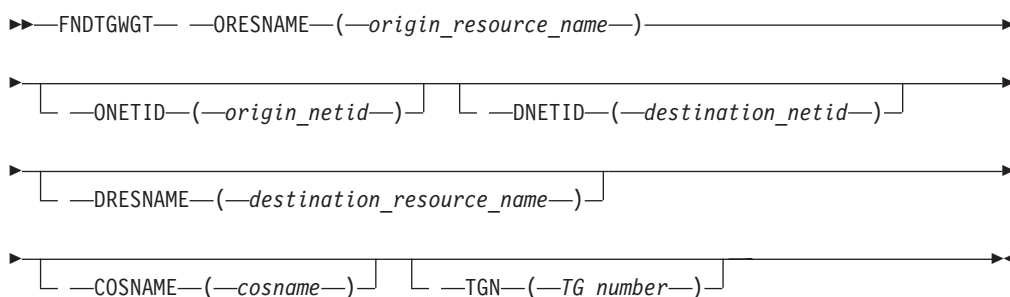
If cosname is not specified, the default cosname is #CONNECT.

TG_number

The TG number must be a decimal number from 0 - 255. If the TG number is not supplied, it displays all TGs.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

FNDTGWGT ORESNAME(SSCP1A)

FNDTGWGT Analysis

Note - Reason will be present only if TG weight is 32767 (infinite).
The infinite weight TG can not be used in route calculation.

Dest CPNAME	TG Addr	TGN	TG Weight	Reason
NETA.SSCP2A	16B5E2D0	21	32767	INOP
NETA.SSCP2A	16B5E850	22	30	

HOST

Use HOST to determine the following items for the VTAM host from which a dump was taken:

- Address space ID (ASID)
- CP name
- CP network address
- Host CDRM name
- Network ID
- Node type
- SSCP network address
- Whether the CDRM supports CDRSCs
- Whether the SSCP is gateway-capable

This information can provide a valuable point of reference for network problem diagnosis.

Use the following syntax as an alternative to the panel interface.

Syntax

▶▶—HOST—▶▶

Equated symbol

Symbol

Description

ATCVT

VTAM ATCVT.

ISTSRTDcdrmname

Derived from the CDRM RDTE for the host. (For example, for cdrmname SSCP1A, the ISTSRTDcdrmname symbol would be ISTSRTDSSCP1A.)

Sample output

HOST

HOST Analysis

```

NetID           NETA
ASID (Hex)      001D
ASID (Dec)      29
Subarea (Hex)   0000000D  Element  0001
Subarea (Dec)   13      Element   1
CDRM Name       A13N
    
```

```

This SSCP is gateway capable
This CDRM supports dynamic CDRSCs
    
```

```

CP Network address 0000000D 0006
CP Name           NETA.A13N
This is a pure network node
    
```

HPRIP

Use HPRIP to format control blocks associated with Enterprise Extender support.

Operands

Line name

The line name should be 1–8 alphanumeric characters. If it contains fewer than eight characters, it is padded on the right with blanks.

You can also specify HPRIP without the operand. If you do, you will get the control blocks associated with all of the lines defined in the Enterprise Extender XCA major node.

Use the following syntax as an alternative to the panel interface.

Syntax

```

  >> HPRIP [—LINEN(line_name)] >>

```

Sample output

HPRIP LINEN(LN1A2A)

```

                                HPR/IP Analysis
AUVT: 02556018
+0000 C1E4E5E3 025869F0 8253E014 00000000 | AUVT...0b..... |
+0010 82538014 00000000 8258CD34 8258C46C | b.....b...b.D% |
+0020 825894A4 83FB56F8 8258A43C 025C8DF0 | b.muc...8b.u..*.0 |
.
.
.
IPNCB: 025869F0
+0000 16400000 0258BA08 03C9F088 00050000 | . ....I0h.... |
+0010 00000001 00000000 00000000 00000000 | ..... |
.
.
.
+0170 0256E340 0256E340 00000001 00000001 | ..T ..T ..... |
+0180 10000400 09F94BF1 F84BF5F2 4BF10000 | ....9.18.52.1.. |
+0190 00000000 83DAA2BC 83DAA2BC 83DAA2BC | ...c.s.c.s.c.s. |
+01A0 83DAA2BC 8258A43C 00000008 0001010E | c.s.b.u..... |
+01B0 00010112 00010112 00000000 E7C3C1F1 | .....XCA1 |
+01C0 C1404040 00000000 00000000 00000000 | A ..... |
+01D0 D7D6D9E3 F1C14040 02555018 204080C0 | PORT1A ..&.. .. |
+01E0 00000000 00004040 2EE42EE3 2EE22EE1 | ..... .U.T.S.. |
+01F0 2EE00000 00000000 00000000 | ..... |
.
.
.
UDATA: 02555018
+0000 01A00000 09123401 E3C3D7C3 E2F64040 | .....TPCS6 |
+0010 2EE00005 82AFC530 00000000 00000000 | ....b.E..... |
.
.
.
AUNCB FOR LN1A2A
AUNCB: 0256E340
+0000 15640061 0258BBB8 03C9F088 00050000 | .../.....I0h.... |
+0010 00000001 00000000 00000000 00000000 | ..... |
+0020 00000000 00000000 00000000 02529010 | ..... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 02556040 00000000 | .....- .... |

```

HPRIP

```
+0050 BA090000 1FF00040 00000000 00000000 | .....0. .... |  
:  
:  
:
```

ISTVABND

ISTVABND determines the following in an MVS dump of a VTAM abend:

- System completion code
- Program interrupt code
- Instruction length code
- Translation exception address
- PSW
- Abnormally ending module name, displacement, PTF level
- Failing instruction
- Registers at time of abend
- VTAM save area chain (forward and backward)
- Symptom string information

IPCS symbols for each register and the PSW address are created. After ISTVABND executes, storage pointed to by the registers and PSW can be accessed by using these IPCS symbols.

The ISTVABND command can also be issued from the panel interface.

Syntax

▶—ISTVABND—▶

Equated symbol

Symbols	Description
R0, REG0	Register 0
R1, REG1	Register 1
R2, REG2	Register 2
R3, REG3	Register 3
R4, REG4	Register 4
R5, REG5	Register 5
R6, REG6	Register 6
R7, REG7	Register 7

R8, REG8
 Register 8
R9, REG9
 Register 9
R10, REG10, RA, REGA
 Register 10
R11, REG11, RB, REGB
 Register 11
R12, REG12, RC, REGC
 Register 12
R13, REG13, RD, REGD
 Register 13
R14, REG14, RE, REGE
 Register 14
R15, REG15, RF, REGF
 Register 15
PSW PSW address
MODULE NAME
 Module that called ISTSSCZZ

Note: If the abend was the result of an ABEND0A9 issued by module ISTSSCZZ, the registers at the time of the call to ISTSSCZZ (rather than the registers when ISTSSCZZ issued the ABEND0A9) are used to create the symbols listed above. Also, the module name, displacement, and PTF level of the module that called ISTSSCZZ are displayed.

Additional information

An abend can occur in SRB mode or in TCB mode.

For an MVS dump of an abend in TCB mode, ISTVABND locates the abnormally ending TCB and the RTM2WA. The completion code (system or user) in the TCB is analyzed and displayed. From the RTM2WA, the program interrupt code (PIC), instruction length code (ILC), and translation exception address (TEA) if valid, are analyzed and displayed.

For an abend in SRB mode, there is no RTM2WA. The PIC, ILC, TEA, registers, and PSW are taken from the SDWA, which is found in the MVS FRR stack.

The PSW address is used to determine the abnormally ending module name, displacement, and PTF level. Register 13 (which usually contains a pointer to the abnormally ending module's save area) is used to trace the save area chain forward and backward (by calling ISTVSAVE) to show module linkage.

Symptom string information is obtained from the variable recording area (VRA) of the SDWA.

Sample output

The following information shows a sample of the output from ISTVABND for an ABEND0C4 in module ISTTRTLR.

ISTVABND

CLIST ISTVABND STARTED AT 09:06:51.

(ISTVABND) THIS DUMP WAS THE RESULT OF AN ABEND IN SRB MODE

SYSTEM COMPLETION CODE = 0C4
PROGRAM INTERRUPT CODE = 0010 INSTRUCTION LENGTH CODE = 0004

PSW AT TIME OF ABEND: 076C2000 82DE96FA
TRANSLATION EXCEPTION ADDRESS = 30580038
THE FAILING INSTRUCTION IS: 43603001
ISTAPCFR-VTAM FRR DUMP

VTAMMAP input data
VTFNDMOD SYMBOL(PSW) NINTERNAL

Module name: ISTRTRLR
Compile date: 92.224
Address entered: 02DE96FA
Module entry point: 02DE9558

Displacement into module: 1A2

First '40'X bytes of module:

```
DATA: 02DE9558
+0000 47F0F014 0FC9E2E3 E3D9E3D3 D940F9F2 | .00..ISTRTRLR 92 |
+0010 4BF2F2F4 90ECD00C 05C018FD 5860F000 | .224.Ö}..{...-0. |
+0020 58D06000 1F99BF97 C40F58A0 C40A187D | .}-..rPpD..µD..'|
+0030 1E791E7A 50706000 5880C3CE 18B714B8 | .~.:&.-...Cö.¼.½ |
```

Storage around address entered:

```
DATA: 02DE96E6
+0000 5840806C 41500002 1E544130 50304190 | . %.&.....&... |
+0010 00021F66 43603001 1E69D200 30013000 | .....-.....K..... |
+0020 42603000 4130504F 1F664360 30011E69 | .-.....&|...-.... |
+0030 D2003001 30004260 30004110 D0681E94 | K.....-.....}..m |
```

REGISTERS AT TIME OF ABEND:

REG0 = 82D9FE00 REG1 = 028F8010 REG2 = 00000000 REG3 = 3058003A
REG4 = B0580008 REG5 = B058000A REG6 = 00000000 REG7 = 00C49D40
REG8 = 00000002 REG9 = 00000002 REGA = 00000008 REGB = 00000110
REGC = 82DE9572 REGD = 02EC9EA0 REGE = 82DE974C REGF = 00000000

SAVE AREA CHAIN (STARTING WITH SAVE AREA AT 02EC9EA0):

ACRT -> SSTM -> ISTD -> CPNQ
CURRENT SAVE AREA = ACRT

SYMPTOM STRING:

AB/S00C4 LVLS/410 RIDS/ISTAPCFR#R PIDS/5695-11701 ADRS/000001A2
RIDS/ISTRTRLR
LVLS/92.224 REGS/0C188

CLIST ISTVABND ENDED AT 09:07:16. RETURN CODE = 0.

ISTVDUMP

ISTVDUMP determines the SDATA options in effect when an MVS dump occurs. The SDATA options determine which MVS storage areas are requested to be dumped when the dump is taken by VTAM (SDUMP) or requested by the operator (console dump). ISTVDUMP can thus help you determine why a specific address is not in an MVS dump.

Note: ISTVDUMP shows what areas were requested for a dump. However, because the area was requested does not guarantee that information is in the dump. If an area is missing from your dump, it can be due to other reasons (for

example, data is lost transferring the dump from the dump data set to tape, or the dump data set is too small, causing a partial dump to be taken).

The ISTVDUMP command can also be issued from the panel interface.

Syntax

▶—ISTVDUMP—▶

Additional information

When you are working with an MVS dump of VTAM, the following information may be useful:

- The PSA must have been dumped to access low-core address hexadecimal 408.
- CSA must have been dumped to access the ATCVT.
- VTAM private storage must have been dumped (RGN parameter specified when the dump is taken) to access most VTAM modules and control blocks.

If neither CSA nor RGN is requested for a dump, ISTVDUMP issues a message. See the sample output in “Sample output” for an example. To resolve most VTAM problems, you must have the VTAM private region and CSA.

When the dump is taken, ISTVDUMP analyzes the RTM recovery termination control table (RTCT) and the SDUMP parameter list (SDUMP) to determine what was requested on the SDATA operand.

For a stand-alone dump obtained by AMDSADMP, the pointer to the RTCT is 0. If you run ISTVDUMP against a stand-alone dump, a message is issued indicating that the RTCT pointer is 0, and the CLIST stops processing. See Table 48 on page 649 to determine the document that contains information on SDATA options, RTCT, and SDUMP.

The following list shows all of the possible settings of the SDATA flags in the SDUMP parameter list. There is no specific indication for extended areas (above 16 MB). When an area is requested (for example, RGN), it is dumped, as is the extended area if present.

Sample output

In this sample, note that CSA and RGN (SDATA option SDURGN) were both requested.

ISTVDUMP:

CLIST ISTVDUMP STARTED AT 13:36:44.

SDATA OPTIONS REQUESTED FOR THIS DUMP:

```
SDUALPSA - DUMP ALL PSA'S IN THE SYSTEM
SDUPSA   - DUMP THE CURRENT PSA
SDUNUC   - DUMP THE NUCLEUS
SDUSQA   - DUMP SQA
SDULSQA  - DUMP LSQA
SDURGN   - DUMP REGION (PRIVATE AREA)
SDULPA   - DUMP ACTIVE LPA MODULE FOR RGN
SDUTRT   - DUMP TRACE TABLE / GTF BUFFERS
SDUCSA   - DUMP CSA
SDUSWA   - DUMP SWA FOR REGION
```

ISTVDUMP

SDUSMDMP - SUMMARY DUMP REQUESTED
SDUALNUC - DUMP ALL NUCLEUS AREAS

CLIST ISTVDUMP ENDED AT 13:36:44. RETURN CODE = 0.

In this sample, RGN and CSA were not requested when the dump was taken.

ISTVDUMP:

CLIST ISTVDUMP STARTED AT 15:32:17.

SDATA OPTIONS REQUESTED FOR THIS DUMP:

SDUALPSA - DUMP ALL PSA'S IN THE SYSTEM
SDUNUC - DUMP THE NUCLEUS
SDUSQA - DUMP SQA
SDULSQA - DUMP LSQA
SDULPA - DUMP ACTIVE LPA MODULE FOR RGN
SDUTRT - DUMP TRACE TABLE / GTF BUFFERS
SDUSWA - DUMP SWA FOR REGION
SDUSMDMP - SUMMARY DUMP REQUESTED

```
*****  
* PRIVATE REGION WAS NOT DUMPED *  
*****  
*****  
* CSA WAS NOT DUMPED *  
*****
```

CLIST ISTVDUMP ENDED AT 15:32:18. RETURN CODE = 0.

SDUMP parameter list

Flag	Description
------	-------------

SDUALPSA

Dump all PSAs in the system.

SDUPSA

Dump the current PSA.

SDUNUC

Dump the nucleus.

SDUSQA

Dump SQA.

SDULSQA

Dump LSQA.

SDURGN

Dump region (private area).

SDULPA

Dump active LPA module for RGN.

SDUTRT

Dump trace table and GTF buffers.

SDUCSA

Dump CSA.

SDUSWA

Dump SWA for region.

SDUSMDMP
Summary dump requested.

SDUNSM DP
Do not dump summary dump.

SDUNSPSA
Do not dump all PSA.

SDUNASQA
Do not dump SQA.

SDUALNUC
Dump all nucleus areas.

ISTVMAP

Use ISTVMAP to determine the starting and ending addresses and area size of the following major MVS storage areas in a dump:

- CSA
- Extended CSA
- Extended FLPA
- Extended maximum possible region
- Extended MLPA
- Extended PLPA
- Extended private region
- Extended read/only nucleus
- Extended read/write nucleus
- Extended SQA
- FLPA
- Low storage
- Maximum possible region
- MLPA
- PLPA
- Private region
- Read/only nucleus
- Read/write nucleus
- SQA

When you cannot find an address in a dump, the starting and ending addresses of major MVS storage areas in the dump will help you determine whether and where that address is in the dump.

Also use ISTVMAP when areas of storage needed to diagnose a VTAM problem do not appear to be in the dump. Knowing which storage area a given address represents and what was dumped can be helpful in determining why a specific storage address is not in a dump. See “ISTVDUMP” on page 244 to determine which storage areas were requested to be dumped when the dump was taken.


The ISTVMAP command can also be issued from the panel interface.

Operands

ASID (*asid*)

The ASID of the address space can be specified to be mapped. ASID may be specified in decimal or hexadecimal format. Enter hexadecimal values in the form X'xx'. ASID should be in the range of 1 to the maximum number of address spaces in the dump.

Syntax



Additional information

The storage in an MVS system is mapped by the CVT virtual storage address extension, the global data area (GDA), and the local data area (LDA). ISTVMAP uses these areas to produce a map of storage that is unique to the dump being processed. The map does not apply to any other dump.

Some of the ranges in the ISTVMAP output will not be complete if the private region (RGN) of the address space being mapped is not accessible in the dump.

If the ASID parameter is not specified, the ISTVMAP CLIST will use the current ASID (that is, the ASID specified on the IPCS SETDEF command). If that ASID cannot be determined, the ASID from the dump header will be used (the ASID that was current when the dump was taken). If that ASID cannot be determined, VTAM's ASID will be used. If the ASID of VTAM cannot be determined, ASID 0001 will be used.

Sample output

ISTVMAP:

CLIST ISTVMAP STARTED AT 17:34:03.

(ISTVMAP) MAP OF MAJOR MVS STORAGE AREAS FOR ASID X'29', JOBNAME VTAMCS

START	END	SIZE	VIRTUAL STORAGE AREA
0AB00000	7FFFFFFF	75500000	EXTENDED MAXIMUM POSSIBLE REGION
0AB00000	7FFFFFFF	75500000	EXTENDED PRIVATE REGION
06A2D000	0AAFFFFFF	040D3000	EXTENDED CSA
06A2C000	06A2CFFF	00001000	EXTENDED MLPA
06A29000	06A2BFFF	00003000	EXTENDED FLPA
02C44000	06A28FFF	03DE5000	EXTENDED PLPA
01AFB000	02C43FFF	01149000	EXTENDED SQA
01774000	01AFAFFF	00387000	EXTENDED READ/WRITE NUCLEUS
01000000	01773C4F	00773C50	EXTENDED READ/ONLY NUCLEUS
-----16M LINE-----			
00FCF000	00FFFFFF	00031000	READ/ONLY NUCLEUS
00FC0000	00FCEE8F	0000EE90	READ/WRITE NUCLEUS
00E42000	00FBFFFF	0017E000	SQA
00BB9000	00E41FFF	00289000	PLPA
00BB8000	00BB8FFF	00001000	FLPA
00BAB000	00BB7FFF	0000D000	MLPA
00800000	00BAAFFF	003AB000	CSA
00006000	007FFFFFF	007FA000	MAXIMUM POSSIBLE REGION
00006000	007FFFFFF	007FA000	PRIVATE REGION

```
00000000 00005FFF 00006000 LOW STORAGE
CLIST ISTVMAP ENDED AT 17:34:43. RETURN CODE = 0.
```

ISTVSAVE

ISTVSAVE follows a VTAM save-area (VWA) chain forward and backward, starting from the specified save-area address.

Using the save-area address that was entered, ISTVSAVE follows the forward save-area chain until it encounters a 0 or a forward chain pointer that is not valid. Then, starting again at the original save-area address that was entered when ISTVSAVE was invoked, ISTVSAVE follows the backward save-area chain until it encounters a 0 or a backward chain pointer that is not valid. If an error occurs during the attempt to access storage, the save-area chain in that direction (forward or backward) is assumed to end.

The VWA eye-catchers are displayed, separated by arrows (->) indicating the order of VTAM module linkage. If ISTVSAVE encounters a save-area with no (or a not valid) eye-catcher, the output for that eye-catcher may contain unprintable characters or periods (for example, SSUW -> SSZZ -> ...). The (...) means that the first word of the save-area does not contain a standard VTAM save-area eye-catcher.

Operands

You can specify an address or symbol pointing to any valid VWA. You must enter at least one of the following items when you invoke ISTVSAVE:

address

The address of a VTAM module save-area should be 1–8 hexadecimal digits.

symbol

A previously equated IPCS symbol that points to a VTAM save-area.

- X** If the current address being displayed points to a VTAM save-area, the IPCS symbol **X** can be used to represent it.

Optionally, you can specify:

ALL

To list the register save-area for each save-area on the chain.

The ISTVSAVE command can also be issued from the panel interface.

Syntax

```
▶▶—ISTVSAVE—| address and symbol | [ —ALL— ] ▶▶
```

address and symbol

```
| [ —address— ] |
| [ —symbol— ] |
```

Additional information

Most VTAM modules use standard-register save-area linkage. The first word of a register save-area is optional, and some VTAM modules store a 4-character identifier there. These identifiers are the VWA eye-catchers, which are displayed by ISTVSAVE. In most cases, VTAM MVS module names have the form ISTxxCxx, and the identifier consists of the 4th, 5th, 7th, and 8th characters of the name of the VTAM module that owns the save-area.

Example: VTAM Module Name = ISTACCRT VWA eye-catcher = C'ACRT'
 VTAM Module Name = ISTSSCTM VWA eye-catcher = C'SSTM'

If the save area does not follow these rules, it should follow the enhanced save-area chaining convention. If it does not follow the rules or the convention, the results of this CLIST are unpredictable. For the enhanced save-area chaining convention, the first 3 characters are IST followed by the module name. The address pointed to by register 13 always points back to the save-area chain. For a description of both methods, see "Using save-area module linkage conventions—Subarea" on page 405.

Sample output

The following information shows a sample of the output from ISTVSAVE for an ABEND0A9 in module ISTSSCZZ.

ISTSSCZZ is the SSABEND macro processor. The purpose of this module is to issue the ABEND0A9. It is necessary to know the caller of ISTSSCZZ to diagnose the ABEND0A9 properly. In this example, the caller of ISTSSCZZ was ISTSSCUW.

The *current save area* is the save-area pointed to by the address entered when ISTVSAVE was invoked (that is, the save-area for module ISTSSCUW is at address 0B2B0480, which is the address where IPCS was positioned when ISTVSAVE was invoked using the symbol X).

ISTVSAVE R13 ALL:

SAVE AREA CHAIN (STARTING WITH SAVE AREA AT 0B2B0480):

ACRT -> SSTM -> SSTP -> SSKT -> SSU3 -> SSUW -> SSZZ ->

CURRENT SAVE AREA = SSUW

The following information shows a sample of the output from ISTVSAVE for an ABEND0C4 in module ISTDECH2. This is a long save-area chain.

ISTVSAVE 981DE88:

SAVE AREA CHAIN (STARTING WITH SAVE AREA AT 0981DE88):

ACRT -> ACRR -> DEST -> DESD -> DEIS -> DEP2 -> DESA -> DESF -> DEQR ->
 DESF -> DESJ -> DEP2 -> DEVP -> DEK2 -> DESG -> DESB -> DESC -> DEIS ->
 DEK2 -> DEVP -> DEG2 -> DEVP -> DER3 -> DEH2 -> . J. -> .00.

CURRENT SAVE AREA = DEH2

The following information shows sample output from ISTVSAVE with R13 specified as the *symbol* operand. R13 represents address 04FB53C8.

ISTVSAVE R13 ALL:

(ISTVSAVE) CLIST WAS INVOKED WITH ADDRESS/SYMBOL 'R13'
 (ISTVSAVE) SYMBOL 'R13' REPRESENTS ADDRESS 04FB53C8

CURRENT SAVEAREA:

```
SAVE AREA FOR SSNP
04FB53C8.                E2E2D5D7 04FB5360 |          SSNP...-|
04FB53D0. 04FB5448 00000000 00000000 00000000 |.....|
04FB53E0 LENGTH(48)==>All bytes contain X'00'|
04FB5410. 84A3968E                |dto.          |
```

SAVEAREA(S) ENCOUNTERED FOLLOWING THE FORWARD SAVEAREA CHAIN:

```
SAVE AREA FOR ....
04FB5448 LENGTH(76)==>All bytes contain X'00'
```

SAVEAREA(S) ENCOUNTERED FOLLOWING THE BACKWARD SAVEAREA CHAIN:

```
SAVE AREA FOR SSTD
04FB5360. E2E2E3C4 04FB52F0 04FB53C8 84A75C4C |SSTD...0...Hdx*<|
04FB5370. FFA395C8 00000000 04A99CD8 04FB5224 |.tnH....z.Q....|
04FB5380. 04AACF30 04A521F8 00000015 04A395C8 |....v.8....tnH|
04FB5390. 04A521F8 04FB5348 04AACF30 04A99CD8 |.v.8.....z.Q|
04FB53A0. 04AACF08 84A759FA 84A75AAE |....dx..dx!|
```

```
SAVE AREA FOR SSTV
04FB52F0. E2E2E3E5 04FB5288 04FB5360 84A76856 |SSTV...h...-dx..|
04FB5300. FFA759E0 00000000 04A99CD8 04FB5224 |.x.\....z.Q....|
04FB5310. 04FB5224 04AACF30 04FB520C 04FB5348 |.....|
04FB5320. 05003EEE 04FB5360 04A51598 04A99CD8 |.....-v.q.z.Q|
04FB5330. 04AACF08 84A766F2 84A767AE |....dx.2dx..|
```

```
SAVE AREA FOR SSTD
04FB5288.                E2E2E3C4 04FB5170 |          SSTD....|
04FB5290. 04FB52F0 84A75D6A FFA766D8 00000000 |...0dx)|.x.Q....|
04FB52A0. 04A99CD8 04FB5224 04AACF30 04FB520C |.z.Q.....|
04FB52B0. 00000015 04AACF30 04FB520C 00000015 |.....|
04FB52C0. 04A51598 04A99CD8 04AACF48 84A759FA |.v.q.z.Q....dx..|
04FB52D0. 00000000                |....|
```

```
SAVE AREA FOR SSTM
04FB5170. E2E2E3D4 04FB5010 04FB5288 84A41A8A |SSTM..&....hdu..|
04FB5180. FFA759E0 04AACF48 04A99CD8 04A51598 |.x.\....z.Q.v.q|
04FB5190. 00000000 04A51598 0000001D 04A51598 |....v.q....v.q|
04FB51A0. 04AACF30 05003EE6 04FB5224 04A99CD8 |.....W.....z.Q|
04FB51B0. 04AACF48 84A41192 84A4127A |....du.kdu.:|
```

```
SAVE AREA FOR ACRT
04FB5010. C1C3D9E3 00000000 04FB5170 8498E214 |ACRT.....dqS.|
04FB5020. FFA41170 04A99CD8 04A99CD8 00C35558 |.u...z.Q.z.Q.C..|
04FB5030. 84A0D988 04FB5010 00000000 84FB5008 |d.Rh..&....d.&;|
04FB5040. 04A99228 04A99228 04A99228 04A99CD8 |.zk..zk..zk..z.Q|
04FB5050. 8498CD9E 0498DD9D 04A99228 |dq...q...zk.|
```

SAVE AREA CHAIN (STARTING WITH SAVE AREA AT 04FB53C8):

```
ACRT - SSTM - SSTD - SSTV - SSTD - SSNP - ....
CURRENT SAVE AREA = SSNP
```

ISTVSLIP

Use ISTVSLIP to display the registers and PSW that were current at the time of an SLIP dump. The registers and PSW are extracted from the SDUMP buffer pointed to by the CVT.

ISTVSLIP

All 16 general registers and the PSW are displayed, along with the module name and displacement that the address portion of the PSW represents (the module in control at the time the SLIP trap occurred). The module name and displacement that register 14 represents (usually the calling module or within the current module) are also displayed. If the dump was not taken as a result of an SLIP trap, a message to that effect is displayed.

IPCS symbols are created for each register and the address portion of the PSW. After ISTVSLIP has executed, storage locations pointed to by these registers (or PSW) can be displayed using these symbols in the IPCS LIST command. For example, **L R4** will display the storage pointed to by register 4 at the time the SLIP trap occurred.

Note: ISTVSLIP does not support stand-alone dumps taken after the SLIP ACTION=WAIT MVS system command is issued.

The ISTVSLIP command can also be issued from the panel interface.

Syntax

▶—ISTVSLIP—▶

Equated symbol

Symbol or symbols	Description
-------------------	-------------

R0, REG0	Register 0
-----------------	------------

R1, REG1	Register 1
-----------------	------------

R2, REG2	Register 2
-----------------	------------

R3, REG3	Register 3
-----------------	------------

R4, REG4	Register 4
-----------------	------------

R5, REG5	Register 5
-----------------	------------

R6, REG6	Register 6
-----------------	------------

R7, REG7	Register 7
-----------------	------------

R8, REG8	Register 8
-----------------	------------

R9, REG9	Register 9
-----------------	------------

R10, REG10, RA, REGA	Register 10
-----------------------------	-------------

R11, REG11, RB, REGB

Register 11

R12, REG12, RC, REGC

Register 12

R13, REG13, RD, REGD

Register 13

R14, REG14, RE, REGE

Register 14

R15, REG15, RF, REGF

Register 15

PSW PSW address

Additional information

The dump data set name is displayed to verify that the correct dump is being processed. The title of the dump is displayed for additional verification that the dump being processed is, in fact, an ISTVSLIP dump.

To determine the module name, displacement, and PTF level for the PSW and register 14, VTFNDMOD is called. The PSW address and register 14 are used as input to VTFNDMOD.

Sample output

ISTVSLIP

CLIST ISTVSLIP STARTED AT 17:42:19.

DUMP DATASET NAME: IPCS.P620527.DUMPA
 TITLE FROM DUMP: SLIP DUMP ID=0001
 THE ADDRESS OF THE SDUMP BUFFER IN THE CVT IS 00C95000.
 PRIMARY ASID AT THE TIME OF ENTRY TO SLIP IS X'000A'

PSW AT ENTRY TO RTM: 00000000 01D03790

REGISTERS WHEN SLIP TRAP MATCHED:

REG0 = 08000000	REG1 = 080A9000	REG2 = 00000012	REG3 = 00000012
REG4 = 007FF158	REG5 = 80C1EFC8	REG6 = 00003000	REG7 = 00000004
REG8 = 0000000B	REG9 = 81E6190C	REGA = 0650CE40	REGB = 81E51A98
REGC = 81D0347A	REGD = 0650CE40	REGE = 01F2A714	REGF = 010E35E0

(ISTVSLIP) PROCESSING OF PSW FOLLOWS:

VTAMMAP input data
 VTFNDMOD ADDR(X'01D03790') NINTERNAL

Module name:	ISTCFF3D
Compile date:	92.318
Address entered:	01D03790
Module entry point:	01D03462

Displacement into module: 32E

First '40'X bytes of module:

DATA: 01D03462	
+0000 47F0F016 10C9E2E3 C3C6C6F3 C44040F9	.00..ISTCFF3D 9
+0010 F24BF3F1 F80005C0 185F1861 187041F0	2.318..{.-./...0
+0020 00005800 CB7E47F0 C0160000 020089F0=0{.....i0
+0030 0008BFFD C0121B11 0A7818A1 18161807	..P.{.....~....

ISTVSLIP

Storage around address entered:

```
DATA: 01D0377C
+0000 89000018 16101823 58F00010 58F0F034 | i.....0...00. |
+0010 58F0F020 05EF47F0 C5445850 04081255 | .00....0E..&... |
+0020 4780C34A 91405488 4780C34A 58205B50 | ..C+j .h..C+..$& |
+0030 12224780 C34A9102 54884710 C34A5850 | ....C+j..h..C+.& |
```

Address is in extended pageable LPA (above the 16M line).

Extended pageable LPA starting address: 01C13000

Extended pageable LPA ending address: 025F8FFF

(ISTVSLIP) PROCESSING OF REG 14 FOLLOWS:

VTAMMAP input data

VTFNDMOD ADDR(X'1F2A714')

Module name: ISTDSCGD
Compile date: 92.256

Address entered: 01F2A714
Module entry point: 01F2A6A0

Displacement into module: 74

First '40'X bytes of module:

```
DATA: 01F2A6A0
+0000 47F0F014 0FC9E2E3 C4E2C3C7 C440F9F2 | .00..ISTDSCGD 92 |
+0010 4BF2F5F6 90ECD00C 18BF41F0 00005800 | .256..}....0.... |
+0020 B95C0700 47F0B02C 00000200 89F00008 | .*...0.....i0.. |
+0030 BFFDB028 1B110A78 18C150D0 C00450C0 | .....A&}{;&{ |
```

Storage around address entered:

```
DATA: 01F2A700
+0000 B93E58A0 ACE45630 B8E858F0 A0005820 | .....U...Y.0.... |
+0010 B8F805EF 980E1028 58201020 12FF4770 | .8..q..... |
+0020 B094D203 2000B92C 58E04000 50E02024 | .mK.....\ .&\.. |
+0030 47F0B098 41F00008 12FF4770 B52041A0 | .0.q.0..... |
```

Address is in extended pageable LPA (above the 16M line).

Extended pageable LPA starting address: 01C13000

Extended pageable LPA ending address: 025F8FFF

(ISTVSLIP) THE REGISTER SAVE AREA CHAIN COULD NOT BE ANALYZED.
REG 13 IS NOT VALID OR STORAGE IS NOT AVAILABLE.

CLIST ISTVSLIP ENDED AT 17:43:30. RETURN CODE = 0.

MNPSC, MNPSD, MNPSF

There are three Multi-Node Persistent Session (MNPS) functions: MNPSD, MNPSC, and MNPSF. They provide two important capabilities:

- Comparing session control blocks from a dump created at the time of a VTAM failure (Z NET,CANCEL or ABEND) with a dump created after the recovery has completed. The **MNPSD** function creates a dump of the MNPS-related control blocks of the MNPS sessions when there is a VTAM failure. The **MNPSC** function compares the control blocks of that dump with the ones in a dump created after the recovery to determine whether session characteristics are the same.
- Formatting MNPS session control blocks using the **MNPSF** function.

MNPSF, MNPSD, MNPSF



Sample output

MNPSD DDNAME(MNPSDUMP)

MNPS Dump Process
00000004 Record(s) were dumped successfully

MNPSF DDNAME(MNPSDUMP) RESOURCE(MAPPC2A1)

MNPS Compare Process
FMCB Extension

Field	Current	Dataset	Field	Current	Dataset
EXLEN	84	84	FMPRO	03	03
NAME	TCPM1011	TCPM1011	NETID	NETA	NETA
BRQS	1	1	BRQR	0	0
BRPS	0	0	BRPR	1	1
SSS	1	1	RNAMS	0	0
:					

ISTFMCB					
Field	Current	Dataset	Field	Current	Dataset
TYPE	03	03	LNGLTH	0286	0286
MXRUI	00000400	00000400	MXRU0	00000F00	00000F00
AVAIL	1	1	OCFLG	0	0
ICFLG	0	0	LU6	0	0
SESTY	0	0	ASPI	1	1
:					

ISTRPNCB					
Field	Current	Dataset	Field	Current	Dataset
REMOTE_CONN_ID	1427431700000076	1427431700000076			
CONN_STATE	32	32	MAX_DATA_SIZ	0000E01D	0000E01D
PATHSWITCH_T	000000F0	000000F0	LOCAL_MNPS	1	1
ALS_CPNETID	NETA	NETA	ALS_CPNAME	SSCP2A	SSCP2A
RELIABLE	1	1	CONN_TYPE	11	11
:					

ISTBSB					
Field	Current	Dataset	Field	Current	Dataset
BSBID	BC	BC			
MAXSL	*00	87	MAXPL	*00	F8
F5SA_MSG	00000000000000050000000000000003	00000000000000050000000000000003			
TYPE	1	1	IND	1	1
:					

ISTSIB					
Field	Current	Dataset	Field	Current	Dataset
CBID	98	98			
PCID	EAABEEC39E9C6EA5	EAABEEC39E9C6EA5			
BDEVN	0000000000000000	0000000000000000			
BCKUP	0	0	BXRFS	0	0
BSCI	1	1	BDLUT	*0	1
:					

End of file reached
00000004 Record(s) were scanned successfully

MNPSF RESOURCE(MAPPC2A1)

MNPS Format Process

FMCBX: 08B12138					
+0000	00000001	001D8402	02000004	08B0B018d.....
+0010	00000001	001B0000	0001001D	08B120A8y
+0020	00000000	08D4D7F8	00000000	00000000MP8.....
+0030	00000000	D3F7F2F1	F1C14040	D5C5E3C1L7211A NETA
+0040	40404040	00000000	00000000	00000000
+0050	00000000	00009080	00000000	00100000
+0060	00000000	0A05C034	EAABEEC3	9E9C6EA1{.....C...>~
+0070	00000000	00000000	D3F7F2F1	F1C14040L7211A

```

+0080 00000000 | .... |
FMCB: 08B0B018
+0000 03008000 08B12138 08B59B90 00000000 | ..... |
+0010 08C68210 00000000 00000000 00000000 | .Fb..... |
+0020 08C67260 00000000 00000000 1FF00018 | .F.-.....0.. |
+0030 00000000 00000000 0101001C 01010021 | ..... |
+0040 00000000 00000000 08C67030 00000000 | .....F..... |
+0050 1D096800 0FF00040 00000000 00000000 | .....0. .... |
+0060 0101001C 01010021 00000000 00000000 | ..... |
+0070 00000000 00000000 00000000 02000004 | ..... |
:
RPNCB: 08B2B800
+0000 FABF001C 09F8B628 08C95048 00050000 | .....8...I&.... |
+0010 0000000B 00000000 00000000 00000000 | ..... |
+0020 00000000 00000000 00000000 00000000 | ..... |
+0030 00000000 00000000 00000000 00000000 | ..... |
+0040 00000000 00000000 090468B4 00000000 | ..... |
+0050 95001200 1FF00040 00000000 00000000 | n....0. .... |
+0060 00000000 00000000 89C584E0 00000000 | .....iEd\.... |
+0070 37010000 0FF00060 00000000 00000000 | .....0.-..... |
:
BSB: 08B14228
+0000 BC0000C0 00800080 00040080 00800001 | ...{..... |
+0010 00000000 00000000 00000000 00000000 | ..... |
+0020 00000000 00000000 00000000 00004000 | ..... |
+0030 00000000 00000003 00000000 00000001 | ..... |
+0040 00000001 001B0000 0000001D 00000001 | ..... |
+0050 001C3C16 D10040C0 00017FFF 00010007 | ....J. {..". |
+0060 00070001 00000000 00000000 00000000 | ..... |
+0070 0000260D 00000000 00000000 00000000 | ..... |
:
SIB: 09FE9300
+0000 9800FC00 00000000 EAABEEC3 9E9C6EA1 | q.....C..>~ |
+0010 00000000 00000000 40404040 40404040 | ..... |
+0020 00000000 00000000 B3942810 BDE39308 | .....m...Tl. |
+0030 09FE9188 09FE9478 00000000 00000000 | ..jh..m..... |
+0040 00000000 09FE9410 09FE93B0 0A05C034 | .....m...l...{. |
+0050 10311000 00000008 09FE9478 00000000 | .....m..... |
+0060 00000000 00000000 A0000000 08B18018 | ..... |
+0070 00000090 00000000 00000000 00000000 | ..... |
:

```

PABSCAN

A request/response unit processing element (RUPE) represents the unit of work VTAM must perform for a given request or response received from the network. VTAM queues RUPes to a process anchor block (PAB). Knowing what type of work is queued to a PAB may be important in resolving storage and performance problems. Use PABSCAN to scan a chain of RUPes queued to a PAB and obtain a summary of the RUPes by RU type.

You may scan all work elements on the PAB, or limit the search to RUPes containing a specific value in one or more of the following fields:

- Destination address field (DAF)
- Origin address field (OAF)
- Request/Response unit (RUPERQD,RUPERSD)

- A user-specified location within the RUPE

For each work element that is selected, the RU is extracted and counted. After all of the selected work elements have been counted, a summary showing the number of work elements containing each RU type is displayed.

Operands

You must specify one address or one symbol to represent the first RUPE in the chain of RUPES to be analyzed.

Address

Enter 1–8 hexadecimal digits in the form X'x...'. for the address of the chain of RUPES to be analyzed. If the address is fewer than 8 digits, it is padded on the left with zeros.

IPCS symbol

Enter an IPCS symbol name that is 1–31 alphanumeric characters. The symbol name represents the beginning of a chain of RUPES. Do not include a period.

Under IPCS, the symbol X represents the address currently being displayed. If the current address is pointing to a chain of RUPES, this symbol may be used.

If you specify no other selection operands, the first 100 RUPES are analyzed and a summary of the RUPES by RU type is displayed.

If you use more than one of the following operands, all of the selection criteria must be met for a RUPE to be selected.

Destination address field

Only RUPES containing this destination address field (RUPEDAF) are eligible for selection. Specify 1–12 hexadecimal digits in the form X'x...'. If the address is fewer than 12 digits, the rightmost digits are compared.

Origin address field

Only RUPES containing this origin address field (RUPEOAF) are eligible for selection. Specify 1–12 hexadecimal digits in the form X'x...'. If the address is fewer than 12 digits, the rightmost digits are compared.

Control op code

Only RUPES containing this CPCB op code are eligible for selection. The control op code must be 1–8 hexadecimal digits in the form X'x...'. If the op code is fewer than 8 digits, it is left-aligned and compared with the leftmost digits in the dump.

Request/response unit

Only RUPES that contain this RU are eligible for selection. The leftmost digits of field RUPERSD are compared (if RUPERSP is nonzero) to the value entered; otherwise, the leftmost digits of RUPERQD (if RUPERQP is nonzero) are compared to the value entered. The length used for the comparison is the length of the value entered.

Detail

The default is N. Specify Y to have the following fields extracted and displayed for each RUPE meeting the selection criteria:

- Position of RUPE on the PAB
- RUPE address
- Origin address field (RUPEOAF)
- Destination address field (RUPEDAF)

- First 4 bytes of request/response unit (if present)
- User-data at a specified displacement (if Displacement and Value are specified)

One line of output per RUPE is produced.

Max

Specify the maximum number of RUPes to be processed. The default for MAX is 100. If MAX is not specified, only the first 100 RUPes on the PAB are analyzed. The maximum value for MAX is 99999. If the maximum number of RUPes are processed and more remain on the PAB, PABSCAN will report the number of unprocessed elements remaining on the PAB.

Displacement

Enter the displacement into the RUPE where Value is to be found. The maximum decimal displacement is 4095 and the maximum hexadecimal displacement is X'FFF'.

Length

Enter a value of 1–8 for the number of bytes you want displayed, starting at the displacement specified in Displacement.

Length must be used with the Displacement operand. Together, they display any portion of a RUPE. The Length operand cannot be used with the Value and Value Type operands.

Note: The following two operands, Value and Value Type, must be used together with the Displacement operand. They allow any field in a RUPE to be checked for a user-specified value. The Value and Value Type operands cannot be used with the Length operand.

Value

Only RUPes containing this data at the displacement specified in Displacement are eligible for selection.


Value may contain character or hexadecimal data of 1–8 bytes in length. Character data should consist of alphanumeric characters. Hexadecimal data should contain an even number of up to 16 hexadecimal digits in the form X'xx...'; otherwise, the high-order half-byte is assumed to be 0.

Binary data can be used to look at a particular bit within a byte. You may specify 1 byte of binary data in the form X'xx'. Only 1 bit within the byte may be selected. Therefore, you can specify only the following hexadecimal values: 01, 02, 04, 08, 10, 20, 40, and 80. A value with more than 1 bit set (for example, 82) will not be processed.

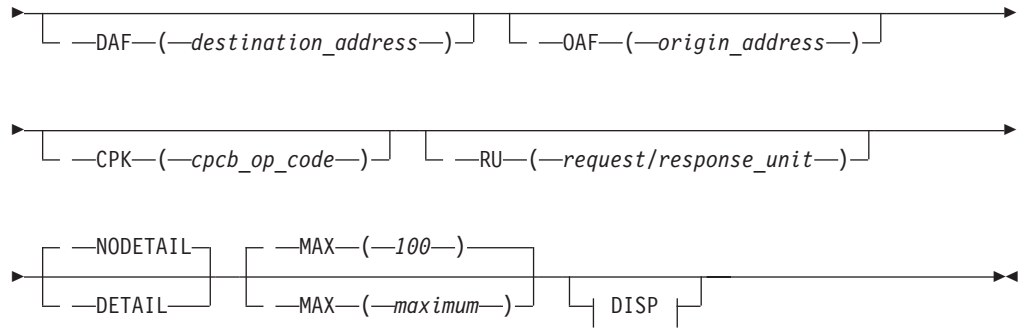
Value type

Enter B for binary, C for character, or X for hexadecimal to indicate the type of data entered for Value.

Syntax



PABSCAN



DISP



VAL



TYPE



LEN



Additional information

To determine whether an RU is a request or a response, PABSCAN first checks field RUPERSP.

If RUPERSP is nonzero, the RU is considered to be a response. The designation RSP, along with the contents of field RUPERSD, are used to represent the response.

If RUPERSP is 0, field RUPERQP is then checked.

If RUPERQP is nonzero, the RU is considered to be a request. The designation REQ, along with the first 4 bytes of the RU from field RUPERQD, are used to represent the request (using 4 bytes allows for the largest of the RU headers and also picks up the format byte for RU headers which are 3 bytes long).

If both RUPERQP and RUPERSP are 0, the designation NORU, along with the contents of field CPCBOPC, are used to represent the RU. The designation NORU notes the fact that no RU (neither request nor response) pointer existed in the RUPE or the length of the RU was 0.

Use the following syntax as an alternative to the panel interface.

Sample output

PABSCAN SYMBOL(X)

ELEM#	RUPEADDR	CPCBOPC	PABSCAN Analysis			USERDATA
			RUPEOAF	RUPEDAF	RU	
1	062E3028	00000000	000000010003	000000010003	**NORU**	
	00000000	REQ occurred	1			
		RUPes left on the chain	0			
		Elements processed	1			

PARTNRLU

Use PARTNRLU to display all partner LUs for an APPC application. PARTNRLU formats and displays the APPCB control block, the COPR control block if present, and the LME.

Operands

APPC application name

The APPC application name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

The APPC application name is required.

Use the following syntax as an alternative to the panel interface.

Syntax

►►—PARTNRLU— —APPLNAME—(—APPC_application_name—)—————►►

Sample output

PARTNRLU APPLNAME(APPCAP05)

PARTNRLU Analysis						
APPCB: 0290E6B8						
APPLUCB..	0291A100	APPTSKID.	02818588	APPACB...	00CB4820	
APPLUN...	APPCAP05					
APPSPTAE	02906530	0290D898	02906620	029065D0	02906580	
DATA: 0290E6B8						
+0000	62C1D7D7	0291A100	02818588	00000000		.APP.j~.aeh....
+0010	00000000	00000000	024100B0	00000000	[....
+0020	31094000	00000010	00000000	00000000	
+0030	0101001B	00000000	00CB4820	0290B088	[h
+0040	00000000	00000000	00000000	00000000	
+0050	00000000	00000000	00000000	00000000	
+0060	C1D7D7C3	C1D7F0F5	00000000	00000000		APPCAP05.....
+0070	00000000	00000000	028030C0	02804028	{. .
+0080	00000000	0292A2FC	00000000	00000000	ks.....
+0090	00000000	0290DA98	00000000	00000000	q.....
+00A0	00000000	00000000	023FA0A8	00000000	μy....
+00B0	36200000	000000A0	02906530	0290D898	μ.....Qq
+00C0	02906620	029065D0	02906580	00000000	}.....
+00D0	00000000	00000000	00000000	00000000	
+00E0	00000000	00000000	00000000	00000000	
+00F0	00000000	00000000	00000000	00000000	
+0100	00000000	00000000			

No session limit negotiations were in progress

Current Partner LU(s) for APPCAP05
 LME: 0290B148
 LMENETID. NETA LMENM.... APPCAP06 LMEFSM... C2

RDTCHECK

RDTCHECK displays the RDTE name, RDTE address, RDTE entry type, RDTE header type, network address, and the current and desired state of an RDTE. In addition, RDTCHECK displays pertinent flag bits from the following control blocks if available:

- Resource definition table application entry (RAP)
- Resource definition table physical unit entry (RCC)
- Resource definition table cross-domain resource manager (RCDRM)
- Resource definition table cross-domain resource entry (RCDRS)
- Resource definition table allocation entry prefix (RCPRE)
- Resource definition table line entry (RLN)
- Resource definition table logical unit entry (RLU)
- Common physical unit prefix (RPU)
- Resource definition table NCP entry (RRN)

Operands

You must specify one address or one symbol.

Address

Enter 1–8 hexadecimal digits in the form X'x...!' for the address of the RDTE to be analyzed. If the address is fewer than 8 digits, it is padded on the left with zeros.

IPCS symbol

Enter 1–31 alphanumeric characters for an IPCS symbol name that has been previously equated with the address of the RDTE to be analyzed. Do not include a period.

Under IPCS, the symbol X represents the address currently being displayed. If the current address is an RDTE, this symbol may be used to refer to it.

Use the following syntax as an alternative to the panel interface.

Syntax

```

  >> RDTCHECK [ —ADDR— (—address—) ] [ —SYMBOL— (—symbol—) ] >>
  
```

Sample output

RDTCHECK ADDR(X'02CE0776')

```

                                RDTCHECK Analysis
RDTE: 02CE077C
RPRNAME.. APPCAP05 RPRENTRY. 55          RPRBITAN. 01000910 01
RPRDEVCH. C06D0000 00800000
DATA: 02CE077C
+0000 C1D7D7C3 C1D7F0F5 80000000 00550200 | APPCAP05..... |
+0010 40040000 0001008D 00060000 02CE0878 | ....._....." |
+0020 02CE07E4 0000015C 0000015C 00000000 | ."U...*...*.... |
+0030 02CE0008 00000000 00000000 05050505 | ."..... |
  
```

```

+0040 00010009 10010010 00000000 00000000 | ..... |
+0050 00000001 00000000 C06D0000 00800000 | .....{|..... |
+0060 00000000 00000000 02CE11A8 00000000 | .....".y.... |
+0070 02CE1020 00000000 00000000 00000000 | ..... |
+0080 00000000 00000000 00000000 00000002 | ..... |
+0090 00000000 00000000 00700033 38E40000 | .....U.. |
+00A0 00000000 00000000 | ..... |

```

```

RPRENTY X'55' indicates an application
RPRHDYX X'02' indicates an application header
Network address      X'00000001008D'
Current state of RDTE X'0505'
Desired state of RDTE X'0505'

```

----- RPRE STATUS BIT FLAGS -----

```

RPRAOPN = 1 Supports LU to LU sessions
RPRDOM  = 1 LU is in this domain
RPRDINUS = 1 This node has been activated at least once
RPRGIST  = 1 Initial status from system definition
RPRDAFAD = 1 RDT added by config services

```

----- RCPRE STATUS BIT FLAGS -----

```

Non-Backup Session Count X'00000002'
Session Limit (Zero Means No Limit) X'0000'
Backup Session Count     X'0000'

```

LOGAPPL

```

RCPRRECD = 1 Record ok
RCPCROSS = 1 Supports cross domain sessions
RCPPRIM  = 1 LU is primary capable
RCPCYMOD = 00 Operator modifiable feature -None
RCPCYSET = 00 SYSDEF defined feature -None
RCPCSM   = 0001 Unstable
RCPSEC   = 1 LU is secondary capable
RCPLVL   = 1 Level of VTAM >= 4.1 for an LU in an NCP segment
RCPUNRCV = 1 Receipt of unrecognized control vector on CINIT supported
RCPSLUSS = 1 Session started is sent by resource when acting as SLU
RCPT21NS = 1 T2.1 nodes and extended BIND supported

```

----- RAP STATUS BIT FLAGS -----

```

RAPASLGI = 1 Application first time logon issued
RAPPARS  = 1 Parsess(yes) was coded
RAPAPASS = 1 CLSDST pass authorized
RAPAACQ  = 1 Acquire authorized
RAPAPPC  = 1 APPC=yes was coded
RAPSRBX  = 1 Schedule exits in SRB mode (OS/VIS only)

```

RDTFULL

Use RDTFULL to display all resource definition table entries (RDTEs) and node control blocks (NCBs) or a selected RDTE.

Operands

RDTE name

The RDTE name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

Use the following syntax as an alternative to the panel interface.

Syntax

```

▶▶ RDTFULL [ —RDTE— (—name—) ] ▶▶

```

Note: Using this function without specifying a name may produce large amounts of data.

Sample output

RDTFULL

```

                                RDTFULL Analysis
QAB: 02955740
+0000 D9C4E340 000D0000 00CC1AB0 02CE0008 | RDT .....".[".. |
+0010 00100000 00700074 | ..... |
RDT: 00CC1AB0
RPRNAME.. VTAMSEG RPRENTRY. 02 RPRBITAN. 02000100 00
RPRDEVCH. 00000000 00000000
DATA: 00CC1AB0
+0000 E5E3C1D4 E2C5C740 80000000 00020200 | VTAMSEG ..... |
+0010 00000000 00000000 00000000 00CC1B18 | .....". |
+0020 00000000 000000B0 00000000 00000000 | .....["..... |
+0030 00000000 00000000 00000000 05050505 | ..... |
+0040 00020001 00000000 00000000 00000000 | ..... |
:
RDTE: 00CC1B60
RPRNAME.. SSCP1A RPRENTRY. 11 RPRBITAN. 00000940 00
RPRDEVCH. C06D0000 00800000
DATA: 00CC1B60
+0000 E2E2C3D7 F1C14040 80000000 00110200 | SSCP1A ..... |
+0010 00000000 00010001 00020006 00CC1C5C | .....".* |
+0020 00CC1BC8 00000210 000000B0 00000000 | ".H.....["... |
+0030 00CC1AB0 00000000 00000000 05050505 | ".["..... |
+0040 00000009 40000100 00000000 00000000 | .... |
:
RDTE: 02CE0CEC
RPRNAME.. APPCAP09 RPRENTRY. 55 RPRBITAN. 09000810 01
RPRDEVCH. C06D0000 00800000
DATA: 02CE0CEC
+0000 C1D7D7C3 C1D7F0F9 80000000 00550200 | APPCAP09..... |
+0010 00000000 00010095 000A0000 02CE0DE8 | .....n.....".Y |
+0020 02CE0D54 0000015C 0000015C 00000000 | ".*.....*.... |
+0030 02CE0008 00000000 00000000 02000200 | ..... |
+0040 00090008 10010010 00000000 00000000 | ..... |
:
+0130 00000000 00000000 00000000 00000000 | ..... |
+0140 00000000 00020001 00010000 00130000 | ..... |
+0150 00000000 00000000 00000000 | ..... |

```

RDTHIER

If an RDTE name is specified, the specified RDTE and all RDTEs below it in the RDTE hierarchy are displayed. If no RDTE name is specified, RDTHIER is identical to RDTFULL.

Operands

RDTE name

The RDTE name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

If the specified RDTE occurs more than once in a dump, the hierarchies for each RDTE are displayed.

Use the following syntax as an alternative to the panel interface.

Syntax

```

►► RDTHIER ───────────────────────────────────────────────────────────────────────────────────────────►►
└── RDTE—(—name—) ─┘

```

Sample output

RDTHIER RDTE(VTAMSEG)

```

RDTHIER Analysis
RDT: 00CC1AB0
RPRNAME.. VTAMSEG  RPRENTRY. 02          RPRBITAN. 02000100 00
RPRDEVCH. 00000000 00000000
DATA: 00CC1AB0
+0000 E5E3C1D4 E2C5C740 80000000 00020200 | VTAMSEG ..... |
+0010 00000000 00000000 00000000 00CC1B18 | .....". |
+0020 00000000 000000B0 00000000 00000000 | .....[..... |
+0030 00000000 00000000 00000000 05050505 | ..... |
+0040 00020001 00000000 00000000 00000000 | ..... |
:
RDTE: 00CC1B60
RPRNAME.. SSCP1A  RPRENTRY. 11          RPRBITAN. 00000940 00
RPRDEVCH. C06D0000 00800000
DATA: 00CC1B60
+0000 E2E2C3D7 F1C14040 80000000 00110200 | SSCP1A ..... |
+0010 00000000 00010001 00020006 00CC1C5C | .....".* |
+0020 00CC1BC8 00000210 000000B0 00000000 | ".H.....[... |
+0030 00CC1AB0 00000000 00000000 05050505 | ".[..... |
+0040 00000009 40000100 00000000 00000000 | .... |
:
RDTE: 02CE0CEC
RPRNAME.. APPCAP09 RPRENTRY. 55          RPRBITAN. 09000810 01
RPRDEVCH. C06D0000 00800000
DATA: 02CE0CEC
+0000 C1D7D7C3 C1D7F0F9 80000000 00550200 | APPCAP09..... |
+0010 00000000 00010095 000A0000 02CE0DE8 | .....n.....".Y |
+0020 02CE0D54 0000015C 0000015C 00000000 | ".*.....*... |
+0030 02CE0008 00000000 00000000 02000200 | ". |
+0040 00090008 10010010 00000000 00000000 | ..... |
:
+0130 00000000 00000000 00000000 00000000 | ..... |
+0140 00000000 00020001 00010000 00130000 | ..... |
+0150 00000000 00000000 00000000 | ..... |

```

RDTSUM

Use RDTSUM to display a summary for all RDTEs or for a selected RDTE.

Operands

RDTE name

The RDTE name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

Syntax

```

►► RDTSUM ───────────────────────────────────────────────────────────────────────────────────────────►►
└── RDTE—(—name—) ─┘

```

Sample output

RDTSUM

```

RDTSUM Analysis
QAB: 02955740
+0000 D9C4E340 000D0000 00CC1AB0 02CE0008 | RDT .....".[".. |
+0010 00100000 00700074 | ..... |
VTAMSEG ADDRESS 00CC1AB0 RPRHDTYP 02 RPRENTRY 02 APPLICATION HDR
RPRDAF 000000000000 RPRCURST 0505 ACTIV
SSCP1A ADDRESS 00CC1B60 RPRHDTYP 02 RPRENTRY 11 CDRM
RPRDAF 000000010001 RPRCURST 0505 ACTIV
ISTATA00 ADDRESS 00CC1D70 RPRHDTYP 02 RPRENTRY 55 APPL
RPRDAF 000000010002 RPRCURST 0200 CONCT
ISTNOP ADDRESS 00CC1ED0 RPRHDTYP 02 RPRENTRY 55 APPL
RPRDAF 000000010003 RPRCURST 0505 ACTIV
ISTPDCLU ADDRESS 00CC2030 RPRHDTYP 02 RPRENTRY 55 APPL
RPRDAF 000000010005 RPRCURST 0505 ACTIV
ISTAPNCP ADDRESS 00CC2190 RPRHDTYP 06 RPRENTRY 11 CDRM
RPRDAF 000000010006 RPRCURST 0505 ACTIV
SSCP1A ADDRESS 00CC23A0 RPRHDTYP 02 RPRENTRY 55 APPL
RPRDAF 000000010008 RPRCURST 0505 ACTIV
ISTPUS ADDRESS 00CC2598 RPRHDTYP 01 RPRENTRY 01 PU T4/5
RPRDAF 000000010000 RPRCURST 0505 ACTIV
ISTGROUP ADDRESS 00CC2740 RPRHDTYP 01 RPRENTRY 30 GROUP
RPRDAF 000000000000 RPRCURST 0505 ACTIV
ISTPDILU ADDRESS 00CC27C8 RPRHDTYP 07 RPRENTRY 07 CDRSC SEGMENT
RPRDAF 000000000000 RPRCURST 0505 ACTIV
ISTADJCP ADDRESS 00CC2870 RPRHDTYP 0F RPRENTRY 0F ADJCP MAJ NODE
RPRDAF 000000000000 RPRCURST 0505 ACTIV
ISTCDRDY ADDRESS 02CBDF58 RPRHDTYP 07 RPRENTRY 07 CDRSC SEGMENT
RPRDAF 000000000000 RPRCURST 0505 ACTIV
ISTDSWMN ADDRESS 02CBEF40 RPRHDTYP 04 RPRENTRY 04 SW SNA MAJ NODE
RPRDAF 000000000000 RPRCURST 0505 ACTIV
:
:

```

ROUTES

Use ROUTES to display explicit route table entries (ERTEs) and virtual route blocks (VRBLKs).

Use the following syntax as an alternative to the panel interface.

Syntax

▶—ROUTES—◀

Sample output

ROUTES

```

ROUTES Analysis
ERTE: 067C7020
ERTPTR... 00000000 ERTERN... 00 ERTFLG... 00 ERTTGN... 01
ERTADJSA. 00000001 ERTDSA... 00000001
DATA: 067C7020
+0000 14280000 00000000 00C70000 00008000 | .....G..... |
+0010 00000000 00000001 00000001 00000001 | ..... |
+0020 00000000 00000000 | ..... |
ERTE: 067C70E0
ERTPTR... 067C70B0 ERTERN... 05 ERTFLG... 00 ERTTGN... 01

```

```

ERTADJSA. 00000004  ERTDSA... 00000002
DATA: 067C70E0
+0000 14280000 067C70B0 05830000 00004000 | .....@.[.c..... |
+0010 00000000 00000001 00000004 00000002 | ..... |
+0020 00000000 00000000 | ..... |

ERTE: 067C70B0
ERTPTR... 067C7080  ERTERN... 02          ERTFLG... 00          ERTTGN... 01
ERTADJSA. 00000004  ERTDSA... 00000002
DATA: 067C70B0
+0000 14280000 067C7080 02830000 00008000 | .....@...c..... |
+0010 00000000 00000001 00000004 00000002 | ..... |
+0020 00000000 00000000 | ..... |
:

VRB: 062FE580
VRBADJSA. 00000001  VRBFXCHN. 00000000  VRBDSTSA. 00000001  VRBVRN... 00
DATA: 062FE580
+0000 05E80000 00000000 00000001 00000000 | .Y..... |
+0010 11280000 00000000 062FEE88 00000000 | .....h... |
+0020 00000000 00000000 00000000 00000000 | ..... |
+0030 00C1C610 00000000 1B081000 00000018 | .AF..... |
+0040 00000000 00000000 00000000 00000000 | ..... |
+0050 00000000 00000000 05000000 062E7100 | ..... |
+0060 00000000 00000000 00360000 00000000 | ..... |
+0070 00040000 00000000 00000000 00000000 | ..... |
+0080 05000100 00000000 00000000 00000000 | ..... |
+0090 00000000 00000000 00000000 00000000 | ..... |
+00A0 00000000 00000000 05000200 00000000 | ..... |
+00B0 00000000 00000000 00000000 00000000 | ..... |
+00C0 00000000 00000000 00000000 00000000 | ..... |
+00D0 00000000 00000001 00000000 00000000 | ..... |
+00E0 00000000 00000001 | ..... |

VRB: 062EE268
VRBADJSA. 00000004  VRBFXCHN. 062EE360  VRBDSTSA. 00000002  VRBVRN... 03
DATA: 062EE268
+0000 05E80300 062EE360 00000004 00000000 | .Y...T-..... |
+0010 11280000 00000000 062FEE88 00000000 | .....h... |
+0020 00000000 00000000 00000000 00000000 | ..... |
+0030 00C1C610 00000000 1B081000 00000018 | .AF..... |
+0040 00000000 00000000 00000000 00000000 | ..... |
+0050 00000000 00000000 01000000 00000000 | ..... |
+0060 00000000 00000000 00000000 00000000 | ..... |
+0070 00000000 0000C18 00000000 00000000 | ..... |
+0080 01000100 00000000 00000000 00000000 | ..... |
+0090 00000000 00000000 00000000 0000C18 | ..... |
+00A0 00000000 00000000 01000200 00000000 | ..... |
+00B0 00000000 00000000 00000000 00000000 | ..... |
+00C0 00000000 0000C18 00000000 00000000 | ..... |
+00D0 03000000 00000002 00000000 00000000 | ..... |
+00E0 00000000 00000001 | ..... |
:

```

RTPINFO

Use RTPINFO to display information about RTP pipes.

RTPINFO displays the following information:

- A specific RTP pipe
- All RTP pipes to a particular destination
- All RTP pipes with an exception condition
- All RTP pipes in the system

Operands

ADDRESS

Enter 1–8 hexadecimal digits in the form X'x.....' for the address used to find the RTP pipe. If the address is fewer than 8 digits, it is padded on the left with zeros.

Tip: Use double quotation marks for the address in IPCS for VERBX:

```
VERBX VTAMMAP 'RTPINFO ADDR(X'00000450')
```

SYMBOL

Enter 1–31 alphanumeric characters for an IPCS symbol name that has been previously equated to a location of the RTP pipe. Do not include a period.

Under IPCS, the symbol X represents the address currently being displayed. If the current address points to a location of the RTP pipe, this symbol X may be used to refer to it.

PUNAME

The PUNAME should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, leftmost characters are compared. PUNAME is the name of the RTP pipe.

CPNAME

The CPNAME should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, leftmost characters are compared. CPNAME is the name of the destination CP of the RTP pipe.

NETID

The NETID should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, leftmost characters are compared. NETID is valid only with the CPNAME option. NETID is the network ID of the destination CP of the RTP pipe.

Guideline: If you specify a CPNAME but do not specify a NETID, the host network ID will be used to form a fully qualified network name.

ALLRTPS

ALLRTPS indicates all RTPs in the system are displayed.

EXCEPTN

EXCEPTN indicates all RTPs with the predefined exception condition in the system are displayed. This is the default option when the RTPINFO command is issued without any option.

COUNT

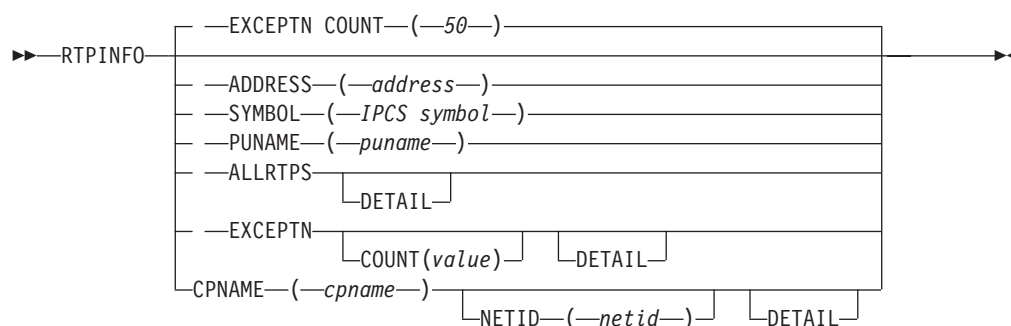
COUNT specifies the minimum number of work elements in one of the RTP queues to display the detail information for the EXCEPTN operand. This operand is valid with the EXCEPTN operand. The valid range is 1–999.

DETAIL

DETAIL indicates that the detail information for each RTP will be displayed. This operand is valid with the CPNAME, ALLRTPS, or EXCEPTN operand. If this operand is not specified, the summary (not detail) is the default value.

Use the following syntax as an alternative to the panel interface.

Syntax



Additional information

- When CPNAME or ALLRTPS option is specified, summary information is displayed for each RTP.
- EXCEPTN is the default option with the count value of 50 work elements. The exception conditions are as follows:
 - RPNCB PABs running with or without queued work elements.
 - The RPN_CONN_STATE is not equal to RPN_CONNECTED (normal state).
 - The RPN_BACKPRESSURE indicator is on.
 - The RPN_NABS value is greater than one.
 - The RPN_NABS_RCV value is greater than one.
 - If the number of work elements in one of the following queues is equal to or greater than the user specified value (or the default value of 50):
 - RPNIBWKQ
 - RPN_PENDING_SENDS_Q
 - RPN_WAIT_FOR_ACK_Q
 - RPN_RCV_MESSAGES_Q
 - RPN_RCV_SEGMENTS_Q
 - RPN_OUTOFSEQ_MSG_Q
 - RPN_PENDING_ALLOC_Q
 - NCBWORKQ

Sample output

RTPINFO PUNAME(CNR00003)

RTPINFO Analysis

Start of CNR00003 Detail information:

```

-----
RTPTB Slot:      000000E6
RPNCB Address:   15147800
RPN_ALS_NAME:    CNR00003
COS Name:        #CONNECT
Destination CPName: SSCP2A
Destination NetID: NETA
RPN_Activation_TOD: 10/16/03 22:01:06.960654
  
```

Connection Information:

```

RPN_CONN_STATE:  X'32' RPN_CONNECTED - Normal State
RPN_LOCAL_CONN_ID: 2E4E0F2A0001000D
RPN_REMOTE_CONN_ID: 2E5289960001000B
RPN_CONN_TYPE:   X'30' - RPN_LULU
                  Not a limited Resource
                  Active End of Pipe
  
```

RTPINFO

RPN_LOCAL_NCB_PTR: 151F2340
Local NCB Type: X'2E'

BackPressure Fields:
RPN_BACK_PRESSURE: OFF
RPN_BP_APPLIED: 0000

BackPressure Reason Counts:
RPN_Backpress_PS_Count: 00
RPN_Backpress_SendQ_Count: 00
RPN_Backpress_Store_Count: 00
RPN_Backpress_Stall_Count: 00

Path Switch Information:
RPN_PSWCH_STATE: OFF
RPN_Cnt_PS_Initiated_Rem: 0000
RPN_Cnt_PS_Initiated_Loc: 0000
RPN_Cnt_PS_Due_To_Failure: 0000
RPN_Cnt_PS_Due_To_PSRETRY: 0000

RPNCB PAB Information:
NCBPCPAB: 00000000 00000000 15EF2F34 00000000
NCBBS PAB: 00000000 00000000 969920C8 00000000
NCBPUPAB: 00000000 00000000 15EF2FA4 00000000

Queue Information:
NCBWORKQ Queue Count: 00000000
RPNIBWKQ Queue Count: 00000000
RPN_Pending_Sends_Q_Cnt: 00000000
RPN_Wait_For_Ack_Q_Cnt: 00000013
RPN_RCV_Messages_Q_Count: 00000000
RPN_OutOfSeq_Msg_Q_Cnt: 00000000
RPN_PENDING_ALLOC_Q_Count: 00000000
RPN_RCV_Segments_NLP_Count: 00000000

Transmission Sequence Numbers:
RPN_NEXT_BYTE_XMIT: 004C4EFF
RPN_LAST_ACK_TRANS: 004C3CD1
RPN_LAST_SREQ_SEQ: 004C4D3C
RPN_LAST_BYTE_RCV: 004C3BF9
RPN_LAST_REXMIT_SEQ: 00000000
RPN_LAST_STATUS_XMIT: 0001
RPN_LAST_ECHO: 0001
RPN_LAST_STATUS_RCV: 0002

RPNCB NAB Information:
RPN_NABS: 00000000
RPN_NABS_RCV: 00000000

RPNCB ARB Information:
ARB_ALLOW_SEND_RATE: 00000C1E
ARB_MAX_SEND_RATE: 00007D00
ARB2_CURRENT_RTT: 00000001
ARB2_SMOOTH_ACTUAL_RATE: 0000034A
ARB_ACCUM_QTIME: 00000B55
ARB2_RCVR_THRESHOLD: 000177AB
ARB2_RCVR_THRESHOLD_MIN: 00004268
ARB2_RCVR_THRESHOLD_MAX: 00009088

End of CNR00003 information -----

Number of RTPs displayed: 1
Number of RTPs found and processed: 1

RTPINFO EXCEPTN COUNT(012)

RTPINFO Analysis

Start of CNR00003 Summary information:

```
-----
RTPTB Slot:      000000E6
RPNCB Address:   15147800
RPN_ALS_NAME:    CNR00003
COS Name:        #CONNECT
Destination CPName: SSCP2A
Destination NetID: NETA
RPN_Activation_TOD: 10/16/03 22:01:06.960654
```

Connection Information:

```
RPN_CONN_STATE:  X'32' RPN_CONNECTED - Normal Stat
RPN_LOCAL_CONN_ID: 2E4E0F2A0001000D
RPN_REMOTE_CONN_ID: 2E5289960001000B
RPN_CONN_TYPE:    X'30' - RPN_LULU
                  Not a limited Resource
                  Active End of Pipe
RPN_LOCAL_NCB_PTR: 151F2340
Local NCB Type:   X'2E'
```

BackPressure Fields:

```
RPN_BACK_PRESSURE: OFF
RPN_BP_APPLIED:    0000
```

BackPressure Reason Counts:

```
RPN_Backpress_PS_Count:  00
RPN_Backpress_SendQ_Count: 00
RPN_Backpress_Store_Count: 00
RPN_Backpress_Stall_Count: 00
```

Path Switch Information:

```
RPN_PSWCH_STATE:      OFF
RPN_Cnt_PS_Initiated_Rem: 0000
RPN_Cnt_PS_Initiated_Loc: 0000
RPN_Cnt_PS_Due_To_Failure: 0000
RPN_Cnt_PS_Due_To_PSRETRY: 0000
```

End of CNR00003 information -----

```
Number of RTPs displayed:      1
Number of RTPs found and processed: 3
```

SES

Use SES to format the RDTE specified by *name*, and all SIBs and RDTEs in session with *name*. The specified *name* can be any session endpoint, such as a logical unit, terminal, or application program. It also formats:

- ACDEBs
- APPCBs
- COPRs
- FMCBs
- FMCBEXTs
- HSICBs
- LUCBs
- NSICBs
- NSSCBs
- RABs
- SABs
- SIBDXs

- SIBIXs
- SIBRXs
- SIBXs

If *name* is not specified, SES formats all SIBs, RDTEs, ACDEBs, APPCBs, COPRs, FMCBs, FMCBEXTs, HSICBs, LUCBs, MPSTs, NSICBs, NSSCBs, PSTs, RABs, SABs, SIBDXs, SIBIXs, SIBRXs, and SIBXs.

Operands

RDTE name

The RDTE name should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks. The specified RDTE can be any session endpoint, such as a logical unit, terminal, or application program.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

```

SES
                                     SES Analysis
MPST: 009E8838
MPSCHAIN. 009E88C0 MPSPSTQ.. 086D2A60
DATA: 009E8838
+0000 D4D7E2E3 00000000 009E88C0 007801F6 | MPST.....h{...6
+0010 00FBC400 086D2A60 00000007 006FB080 | ..D._-.....?..
+0020 C1E4E3C8 88DF92DA 00FD7976 01F60002 | AUTHh.k...~.6..
+0030 00000000 000001F6 00000000 086D2A60 | .....6.....-
+0040 00000C60 00000000 02B70001 009E8838 | .....h.
+0050 00000C60 80000000 00000C60 88DF8930 | .....-h.i.
+0060 00000000 00000000 00000000 00000000 | .....
+0070 000007D8 00000003 | .....Q....
PST: 086D2A60
+0000 61000480 00000000 086D2A60 009E8838 | /....._-..h.
+0010 00000000 00000000 00000000 00000000 | .....
+0020 00000000 00000087 006F8FF0 00000000 | .....g.?..0....
+0030 00000000 00000000 009F8654 00000000 | .....f.....
+0040 01011000 07F00030 00000000 00000000 | .....0.....
+0050 00000000 00000000 009F8658 00000000 | .....f.....
:
ACDEB: 08B9BBE8
ACDTCB... 006FDE48 ACDCHN... 00000000 ACDRDTE.. 09FA1100
DATA: 08B9BBE8
+0000 0F480000 00000000 086D2A60 00000000 | ....._-.....
+0010 00000000 00000000 009F83A0 00000000 | .....c.....
+0020 2D010000 0FF00010 00000000 00000000 | .....0.....
+0030 006FA108 08B9BA98 00000000 00000007 | .?^.....q.....
+0040 00000000 00000000 00000000 0870B188 | .....h
+0050 006FDE48 08B9BA98 00000000 00000000 | .?.....q.....
+0060 00000000 00000000 08C3B1A8 00000000 | .....C.y.....
+0070 08B9BD38 00000000 00000002 00000000 | .....
:
LUCB: 08C3B1A8
+0000 52780076 00000000 086D2A60 00000000 | ....._-.....
+0010 00000000 00000000 08C3B218 00010000 | .....C.....
+0020 00010000 00010000 08000000 09FA1100 | .....
+0030 08B72258 08B722E8 08B9BBE8 00000000 | .....Y...Y....
+0040 00000004 08B72258 00000000 00000000 | .....
+0050 00000000 00000000 009F8398 00000000 | .....cq....
+0060 19011000 0FF00050 00000000 00000000 | .....0.&.....
FMCB: 08B72258
+0000 00000001 00388402 04000002 08B73748 | .....d.....
+0010 00010001 00760000 00010038 08B722E8 | .....Y
+0020 00000000 08B735D8 00000000 00000000 | .....Q.....
+0030 20000000 D3F7F2F0 F1C14040 D5C5E3C1 | ....L7201A NETA
+0040 40404040 00000000 00000000 00000000 | .....
+0050 00000000 00009080 00000000 00000000 | .....
  
```

```

+0060 00000000 00000000 EAABEEC3 946C4B7A | .....Cm%.: |
+0070 00000000 00000000 D3F7F2F0 F1C14040 | .....L7201A |
+0080 00000000 | ..... |
FMCB: 08B73748
+0000 03008040 08B72258 086D2A60 00000000 | .. _-.... |
+0010 08B7B210 00000000 00000000 00000000 | ..... |
+0020 08B7A260 00000000 1C016200 0FF00018 | ..S-.....0.. |
+0030 00000000 00000000 01FF0000 01FF0000 | ..... |
+0040 00000000 00000000 08B7A030 00000000 | ..... |
+0050 1D096200 0FF00040 00000000 00000000 | .....0. .... |
+0060 01FF0000 01FF0000 08B735D8 00000000 | .....Q.... |
+0070 00000000 00000000 00000000 04000002 | ..... |

```

```

:
:
APPCB 08B9B7F8
APPLUCB 08B5D110 APPTSID 08C965B0 APPACB 009E6ED8 APPLUN APPCIA02
APPSPTAE 08B9ABD0 08B9AA68 08B9A888 08B9AAE0 08B9AB58
000000 62C1D7D7 08B5D110 08C965B0 00000000 00000000 00000000 08E7EA88 00000000 *.APP..J..I.....X.....*
000020 31094000 07F00010 00000000 00000000 0101001C 00000000 009E6ED8 08B78078 *...0.....Q.....*
000040 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 *.....*
000060 C1D7D7C3 F1C1F0F2 00000000 00000000 00000000 00000000 00000000 *APPCIA02.....*
000080 00000000 00000000 00000000 00000000 00000000 086DF938 00000000 00000000 *.....9.....*
0000A0 00000000 00000000 08EC86E0 00000000 36200000 07F000A0 08B9ABD0 08B9AA68 *.....0.....*
0000C0 08B9A888 08B9AAE0 08B9AB58 00000000 00000000 00000000 00000000 00000000 *.....*
0000E0 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00010000 *.....*
000100 00000000 00000000 00000000 00000000 00000000 00000000 00000077 00000077 *.....*
000120 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 *.....*
:
:

```

```

FMCBX: 08B725B8
+0000 00010001 00F48403 01000002 086BE218 | ....4d.....S. |
+0010 00000001 00490001 000100F4 08B72408 | .....4.... |
+0020 08B72918 0868E018 08B74438 00000000 | ..... \..... |
+0030 20000000 E3C3D7D4 F1F0F1F1 D5C5E3C1 | .....TCPM1011NETA |
+0040 40404040 00000000 00000000 00000000 | ..... |
+0050 00000000 00009080 00000000 00100000 | ..... |
+0060 00000000 09F96014 EAABEEC3 946C4B7E | .....9-....Cm%= |
+0070 00000000 00000000 E3C3D7D4 F1F0F1F1 | .....TCPM1011 |
+0080 00000000 | ..... |

```

```

FMCB: 086BE218
+0000 03C08000 08B725B8 08C965B0 00000000 | .{.....I..... |
+0010 08B7B210 00000000 00000000 00000000 | ..... |
+0020 08B7A260 00000000 1C016200 0FF00018 | ..S-.....0.. |
+0030 00000000 00000000 0101001C 00000000 | ..... |
+0040 00000000 00000000 08B7A030 00000000 | ..... |
+0050 1D096200 0FF00040 00000000 00000000 | .....0. .... |
+0060 0101001C 00000000 00000000 00000000 | ..... |
+0070 00000000 00000000 00000000 01000002 | ..... |

```

```

:
:
HSICB 086B5258
HSISENSE 00000000 HSICONID 00000000 HSIBIUIN 0B908111 HSISENDQ 00000000
HSIHLDP5 00000000 HSIFMHST 00000000 HSIPACQ 00000000
000000 62C8E2C9 00000000 00000000 00000000 00000000 00000000 0B908111 1580800E *.HSI.....*
000020 00000000 00000000 00000000 00000000 00000000 00000000 00000000 800E0000 *.....*
000040 800E0000 000E0000 00804050 00000000 00050000 00002000 00000000 00000000 *.....*
000060 000E0003 00000000 00000000 00000000 00000000 00000000 00000000 00000000 *.....*
000080 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 *.....*
:
:

```

```

QAB: 08B57060
+0000 D9C4E340 00100000 08B57948 09F80280 | RDT .....8.. |
+0010 00100000 00700074 | ..... |

```

```

RDT: 08B57948
RPRNAME.. VTAMSEG RPRENTRY. 02 RPRBITAN. 02000100 00 RPRDEVCH. 00000000 00000000
DATA: 08B57948
```

```

+0000 E5E3C1D4 E2C5C740 80000000 00020200 | VTAMSEG ..... |
+0010 00000000 00000000 00000000 08B579B0 | .....~ |
+0020 00000000 000000B0 00000000 00000000 | ..... |
+0030 00000000 00000000 00000000 05050505 | ..... |
+0040 00020001 00000000 00000000 00000000 | ..... |
+0050 00000000 00000000 00000000 00000000 | ..... |
+0060 00000000 00000000 20000000 00000000 | ..... |
+0070 08B58418 08B57060 00000000 00000000 | ..d-..... |
+0080 00000000 00000000 00000000 00000000 | ..... |
+0090 00000000 00000000 00000000 00000000 | ..... |
+00A0 00000000 00000000 | ..... |

```

```

RDT: 08B579F8
RPRNAME.. SSCP1A RPRENTRY. 40 RPRBITAN. 00000940 00 RPRDEVCH. C0600000 00800000
DATA: 08B579F8
```

```

+0000 E2E2C3D7 F1C14040 80000000 00400200 | SSCP1A ..... |
+0010 00000000 00010001 00000008 08B57AA8 | .....:y |
+0020 08B57A60 000001D8 000000B0 00000000 | ..?-...Q..... |
+0030 08B57948 00000000 00000000 05050505 | ..... |
+0040 00000009 40000100 00000000 00000000 | ..... |
+0050 00000000 00000000 C06D0000 00800000 | .....{..... |
+0060 00000000 00000000 00000000 00000000 | ..... |
+0070 00000000 00000000 00000000 00000000 | ..... |
+0080 3F013F01 00040000 00080000 00000000 | ..... |
+0090 00000000 00000000 00000000 00000000 | ..... |

```

```

:
:
SIB: 0A07CA58
SIBFSMIN. FC SIBFSMTM. 00 SIBFSENS. 00000000 SIBBPRIQ. 0A07C300 SIBBSEQ. 00000000 SIBTTMFL. 00
SIBTREAS. 00 SIBTSESE. 00
SIBPCID = EAABEEC3946C4B88 QUALIFIER = NETA.SSCP1A

```

```

DATA: 0A07CA58
+0000 9800FC00 00000000 EAABEEC3 946C4B88 | q.....Cm%.h
+0010 40404040 40404040 C9D5E3C5 D9C1C3E3 | INTERACT
+0020 C9D5E3C5 D9C1C3E3 B342A6F4 6489F609 | INTERACT..w4.i6.
+0030 0A07C768 00000000 0A07C300 00000000 | ..G.....C.....
+0040 00000000 0A07CB08 0A07CB68 09F96014 | .....9-
+0050 70311000 00060408 00000000 00000000 | .....
+0060 00000000 00000000 A0000000 00000000 | .....
+0070 80000000 00000000 00000000 00000000 | .....
+0080 09CAE2B0 00000000 00000000 00000000 | ..S.....
+0090 00000000 00000000 00000000 00000000 | .....
+00A0 00000000 00000000 00000000 00000000 | .....
      PLU (OLU) RESOURCE
SIBRX: 0A07CB08
SIBRADJN. .... SIBRALNM. APPC1A02 SIBRNID.. NETA SIBRNETA. 00000001 0049
SIBRNETC = 10 = SIBRNTS - Same domain
DATA: 0A07CB08
+0000 00000000 00000000 C1D7D7C3 F1C1F0F2 | .....APPC1A02
+0010 D5C5E3C1 40404040 00000000 09F93BD0 | NETA .....9.}
+0020 00000000 001073C0 78048100 00000004 | .....{..a.....
+0030 01000007 00000000 00000000 00000000 | .....
+0040 00000000 00010049 00000000 00000000 | .....
+0050 00000000 00000000 00000000 00000000 | .....
      SLU (DLU) RESOURCE
SIBRX: 0A07CB68
SIBRADJN. .... SIBRALNM. AA2LUA1 SIBRNID.. NETA SIBRNETA. 00000003 00E8
SIBRNETC = 10 = SIBRNTS - Same domain
DATA: 0A07CB68
+0000 00000000 00000000 C1C1F2D3 E4C1F140 | .....AA2LUA1
+0010 D5C5E3C1 40404040 00000000 09F803B8 | NETA .....8..
+0020 02020400 001073C0 68008122 10000000 | .....{..a.....
+0030 00000000 7BC9D5E3 C5D94040 02800000 | ...#INTER ...
+0040 00000000 000300E8 0A0A70E4 0A07C768 | .....U..G.
+0050 00000000 09CAE278 00000000 00000000 | .....S.....
:
      DLUR SAW DATA
DATA: 09CAE2B0
+0000 09CAE160 02000000 7BC9D5E3 C5D94040 | .....#INTER
+0010 0A09B010 | ....
SIB: 0A07C768
SIBFMIN. FC SIBFSMTM. 00 SIBFSENS. 00000000 SIBBPRIQ. 0A07C5F0 SIBBSEQ. 0A07C5F0 SIBTTMFL. 00
SIBTREAS. 00 SIBTSESE. 00
SIBPCID = EAABEEC3946C4B83 QUALIFIER = NETA.SSCP1A
DATA: 0A07C768
+0000 9800FC00 00000000 EAABEEC3 946C4B83 | q.....Cm%.c
+0010 C9E2E3E5 E3C3D6E2 C3D7E2E5 D9D4C7D9 | ISTVTCOSCPSVRMGR
+0020 C3D7E2E5 D9D4C7D9 B342A6EA 9A0C8403 | CPSVRMGR..w..d.
+0030 0A07C5F0 0A07CA58 0A07C5F0 0A07C5F0 | ..E0.....E0..E0
+0040 00000000 0A07C818 0A07C878 09F96014 | .....H...H..9-
+0050 50311000 00060408 00000000 00000000 | &.....
+0060 00000000 00000000 A0000000 00000000 | .....
+0070 80000000 00000000 00000000 00000000 | .....
+0080 00000000 00000000 00000000 00000000 | .....
+0090 00000000 00000000 00000000 00000000 | .....
+00A0 00000000 00000000 00000000 00000000 | .....
      PLU (OLU) RESOURCE
SIBRX: 0A07C818
SIBRADJN. .... SIBRALNM. SSCP1A SIBRNID.. NETA SIBRNETA. 00000001 0075
SIBRNETC = 20 = SIBRNTPE - Endpoint
DATA: 0A07C818
+0000 00000000 00000000 E2E2C3D7 F1C14040 | .....SSCP1A
+0010 D5C5E3C1 40404040 00000000 08B58208 | NETA .....b.
+0020 00000000 002073E0 78008100 00000000 | .....\.a.....
+0030 01000005 00000000 00000000 00000000 | .....
+0040 00000000 00010075 00000000 00000000 | .....
+0050 00000000 00000000 00000000 00000000 | .....
      SLU (DLU) RESOURCE
SIBRX: 0A07C878
SIBRADJN. ISTAPNCP SIBRALNM. NNCPA2 SIBRNID.. NETA SIBRNETA. 00000003 00E9
SIBRNETC = 30 = SIBRNTXD - Cross domain
DATA: 0A07C878
+0000 C9E2E3C1 D7D5C3D7 D5D5C3D7 C1F24040 | ISTAPNCPNNCPA2
+0010 D5C5E3C1 40404040 00000000 0A0B5190 | NETA .....
+0020 02020200 003063C0 68808122 10000000 | .....{..a.....
+0030 00000000 E2D5C1E2 E5C3D4C7 07000000 | ...SNASVCMG....
+0040 00000000 000300E9 0A0A70E4 0A07C5F0 | .....Z...U..E0
+0050 0A07CA58 09CAE1D0 00000000 00000000 | .....}.....

```

:

SIBCHECK

SIBCHECK analyzes important fields and relevant status flags in an SIB and related control blocks. The following control blocks are analyzed:

- DLU cross-network extension (SIBX)
- OLU cross-network extension (SIBX)
- PLU resource extension (SIBRX)
- SIB base
- SIB initiation extension (SIBIX)
- SLU resource extension (SIBRX)
- SIB termination extension

SIBCHECK determines:

- Alias resource names and network IDs
- Configuration (cross-domain, cross-network, back-to-back, and so on)
- Destination logical unit
- GWNCP names
- Initiating logical unit
- Network addresses
- Originating logical unit
- Primary logical unit
- RDTE address
- Real resource names and network IDs
- Resource type (APPL, LU, CDRSC, and so on)
- Secondary logical unit

For each status bit in the SIB and related control blocks, the bit name, its value, and its meaning (from) are listed.

Fields in the SIB that contain addresses are checked (such as SIBTV35P, the CV35 pointer, or SIBTNOTP, the pointer to NOTIFY RU). If these fields contain a nonzero address, the address and description of the field are also displayed.

Note: The SIBBTIME value is displayed in the format of coordinated universal time (formerly known as Greenwich Mean Time).

Operands

You must specify one address or one symbol.

Address

Enter 1–8 hexadecimal digits in the form X'x...'¹ for the address of the SIB to be analyzed. If the address is fewer than 8 digits, it is padded on the left with zeros.

IPCS symbol

Enter 1–31 alphanumeric characters for an IPCS symbol name that has been previously equated with the address of the SIB to be analyzed. Do not include a period.

SPANC

SPANC analyzes any or all of the VTAM storage pool anchors (SPANCS). If you use no operands, the number of pages in use, the page size, and where the storage is allocated (common or private) for every SPANC pool are displayed. Options are available to:

- Designate a specific SPANC pool to be analyzed
- Limit the output to CSA or PRIVATE SPANCs
- Determine the number of FBQEs on each page
- Determine the size of each FBQE on each page
- Display a sample of storage from each page
- Determine the page addresses associated with pools
- Process data in a specific pool through the use of an exit

Operands

Pool

Specify the name of a specific SPANC pool to be analyzed. If the pool operand is not used, all SPANCs are processed. If a pool name other than one from the list of valid pool names is specified, no output will be produced.

Note:

1. Pools named 'AVAIL' are not valid and are used only as placeholders. They are displayed in the event storage overlays occur.
2. The FBQE Count, FBQE List, Process, and Exit operands are mutually exclusive; use only one of them.

Pool type

Specifies to format ALL, CSA, or PRIVATE area SPANCs. The default value is ALL. This is valid only when POOL value ALL is specified.

FBQE count

Specify Y to have the number of FBQEs on each page of the selected pool (or all pools if no pool was selected) listed. The FBQE contains the length of the free storage it describes. Use this option for performance or storage fragmentation problems. Long chains of FBQEs can cause VTAM performance problems.

FBQE list

Specify Y to have each FBQE on each page of the selected pool (or all pools if no pool was selected) listed. The FBQE contains the length of the free storage it describes. Use this option for storage fragmentation problems.

Length

Specify the number of bytes of storage you want displayed from the beginning of each page of the selected pool (or all pools if no pool was selected). Any hexadecimal number from X'001' to X'FF8' or any decimal number from 1 to 4088 may be specified. Use this option to get a sample of storage from each page of a specific SPANC pool.

Process

Specify **Map** to display the address of each page that is associated with the selected pool (or all pools if no pool was selected). Use this option with VSMDATA to determine the SPANC pages mapped by each MVS subpool.

Exit

Use Exit to have one of the four exit functions process information on each page of selected SPANC pools. Specify exit FMCB, RU, RUPE, or SIB.

- The FMCB exit searches SPANC pools FMCB, PLUSFMCB, or SSCPFMCB for FMCBs and formats those found.
- The RU exit searches SPANC pools UTILCSAS, UTILCSAL, UTILPVTL, or UTILPVTS for all RUs or a specific RU on a page of storage and displays the address and data for those found.

Note: These pools may contain data that is not an RU. To locate a specific RU, specify the actual RU in the Value field, a Type of X, and a displacement of X'06'.

- The RUPE exit searches SPANC pools RUPECOMM or RUPEPRIV for all RUPEs in the pool. SPANC displays the RUPE address, CPCBOPC, RUPEOAF, RUPEDAF, and RU data for those found.
- The SIB exit searches SPANC pool SIB for all SIBs in the pool. SPANC displays the SIB address, FSMs, sensecode, PLU NetID, PLU name, SLU NetID, SLU name, and procedure correlation identifier (PCID) for SIBs that are found.

Note: For all EXIT routines, an address followed by an asterisk (*) indicates that the buffer pool is allocated.

Note: The following three operands, Displacement, Value, and Value Type, must be used together with the Exit operand. The Exit operand may be used alone.

Displacement

Enter the displacement into the data portion of a page where Value is to be found. The maximum decimal displacement is 4095, and the maximum hexadecimal displacement is X'FFF'.

Value

Enter a character, hex, or binary value to be searched for at the displacement specified by Displacement.

Value may contain character or hexadecimal data of 1–8 bytes in length. Character data should consist of alphanumeric characters. Hexadecimal data should contain an even number of hexadecimal digits in the form X'xx..', otherwise, the high-order half-byte is assumed to be 0.

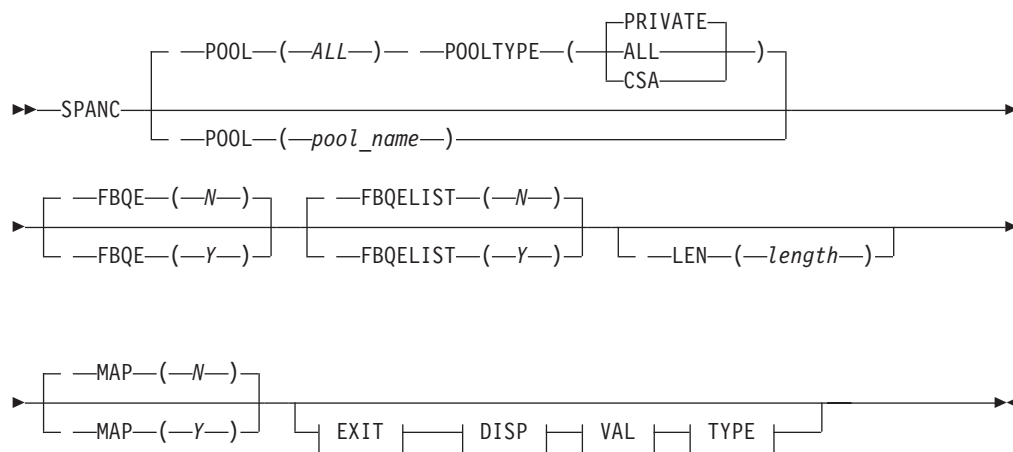
Binary data can be used to look at a particular bit within a byte. You may specify 1 byte of binary data in the form X'xx'. Only 1 bit within the byte may be selected. Therefore, you can specify only the following hexadecimal values: 01, 02, 04, 08, 10, 20, 40, and 80. A value with more than 1 bit set (for example, 82) will not be accepted.

Value type

Enter B for binary, C for character, or X for hexadecimal to indicate the type of data entered for Value.

Use the following syntax as an alternative to the panel interface.

Syntax



EXIT

|— EXIT—(—name—)|

DISP

|— DISP—(—displcmnt—)|

VAL

|— VAL—(—value—)|

TYPE

|— TYPE—(—data_type—)|

Sample output

SPANC

```

                                SPANC Analysis
Page addresses for pool: RUPEPRIV
08D18008
Pages in use for pool RUPEPRIV =      1    Page size = 00001000  PRIVATE
Page addresses for pool: RUPECOMM
080D4010
Pages in use for pool RUPECOMM =      1    Page size = 00001000  COMMON
Page addresses for pool: SIB
091ED008
Pages in use for pool SIB      =      1    Page size = 00001000  PRIVATE
Page addresses for pool: SSCPFMCB
08D65008  0915E008
Pages in use for pool SSCPFMCB =      2    Page size = 00001000  PRIVATE
Page addresses for pool: NQDAT
09123008
Pages in use for pool NQDAT    =      1    Page size = 00001000  PRIVATE
Page addresses for pool: EPTDVT
08388010  08389010
Pages in use for pool EPTDVT   =      2    Page size = 00001000  COMMON
    
```



```

SPANC POOL(RUPEPRIV) EXIT(RUPE)
          SPANC Analysis
RUPE ADDR Op code   RUPEOAF   RUPEDAF   RU Data
-----
02C53020* 08810620 000000010001 000000010010 REQ=8106200302D5376DF4EA
02C530C0* 0F310000 00000001000F 000000010012 REQ=31001307B0B050B30080
02C53160* 0F310000 00000001000F 000000010012 REQ=31001307B0B050B30080
02C53200* 04000000 000000010003 000000010003 REQ=C4C9E2D7D3C1E840C9C4
02C532A0 04000000 000000010003 000000010003
02C53340* 0B310000 00000001000F 000000010012 REQ=FF310281A02801880002
Matches found in exit = 6
Pages in use for pool RUPEPRIV = 1 Page size = 00010000 PRIVATE

```

SRTFIND

Use SRTFIND to locate a symbol resolution table entry (SRTE) in a dump.

Note: An attempt is always made to translate a symbol or a string regardless of the quality of the input data stream. The translation may produce unexpected results such as dots, random letters, or other combinations of symbols.

Operands

SRT name

The SRT name is the symbolic name of a symbol resolution table (SRT) entry and can be entered as alphanumeric characters or hexadecimal digits.

- If alphanumeric characters are used, enter 1–8 characters in the form `cccc`. If fewer than 8 characters are entered, the name is padded on the right with blanks, and the tool will search only for the characters entered.

For example, if `APPL1` is entered, and `APPL1`, `APPL1A`, and `APPL1B` exist, the tool will find only `APPL1`.

- If hexadecimal digits are used, enter an even number of digits 1 - 16 in the form `X'xxxx'`. If fewer than 16 digits are entered, the name is padded on the right with blanks. If an odd number of digits is entered, the name is padded on the left with a 0.

For example, if `X'00000010001'` is entered, the tool will search for `X'0000000100014040'`.

The SRT name is required.

NetID

The NetID name representing the network ID of another network outside the host network where the resource resides should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

Type

Using the name or hexadecimal value, enter the type of SRTE for which you are searching. The default is `RDTE`. Enter hexadecimal values in the form of `X'xx'`.

Chain

Use **Display SRT Chain** to display all SRTEs on the chain, regardless of other search criteria. This option may help locate an SRTE whose storage has been corrupted. The default displays only the SRTEs that match all specified search criteria.

Note: **Chain** overrides the setting for **Process**.

SRTFIND

Format

Use **Format** to have selected data formatted using the SRT control block. **Noformat**, the default, displays the SRT's name, address, and type.

Process

Use **All** or **First** to find the SRTs that match the search criteria. The default is **All**. **First** displays only the first SRTE that matches the search criteria.

Equated symbol

Symbol

Description

ISTSRTsrtname


SRTE

ISTSRTDsrtname

SRTDATA (if present)

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

SRTFIND SRTNAME(SSCP1A) FORMAT

SRTFIND Analysis

SRTE SSCP1A was found at address X'80BF02C8' with type RDTE

```
SRT: 80BF02C8
  SRTSYMM. SSCP1A  SRTSRTE.. 82DA0100  SRTTYPE.. 00
SRTSRTE: 80BF02D0
+0000 SPECE.... 82
DATA: 80BF02C8
+0000 E2E2C3D7 F1C14040 82DA0100 00550200 | SSCP1A b..... |
+0010 00040000 | .... |
```

STORAGE

Use **STORAGE** to format BPCBs, BPDYs, PXBs, SPANCs, and SPTAEs.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

STORAGE

```

                                STORAGE Analysis
:
:
BPD
  DATA: 02952000
    +0000 000C000C 00000000 02952508 00000000 | .....n..... |
    +0010 00000000 00002000 7FFFFFFF 02952390 | .....".n.. |
    +0020 029521A8 00CC41F8 00000000 000003E8 | .n.y.".8.....Y |
    +0030 013400FD 02952054 0281E000 02957000 | .....a\..n.. |
    +0040 00000000 00000000 00000000 00000000 | ..... |
:
:
Buffer pool ID SMS1
  BPCB: 02952390
    BPCBRPHA. 00000000 BPCBRPHB. 00000000 BPCBRPH1. 00000000
    BPCBRPH2. 00000000 BPCBAVNO. 00000000
  DATA: 02952390
    +0000 00000000 00000000 600000E7 02953FF8 | .....-..X.n.8 |
    +0010 00000000 00000000 00000000 00000000 | ..... |
    +0020 00000000 00000000 02952000 00000000 | .....n..... |
    +0030 00000000 00000000 00000000 00000000 | ..... |
    +0040 00000000 00000000 00000000 00000000 | ..... |
:
:
SPANC 02957204
POOLNUM 0000          POOLNAME RUPEPRIV ASSOCID N/A
  DATA: 02957204
    +0000 D9E4D7C5 D7D9C9E5 000C0002 00000000 | RUPEPRIV..... |
    +0010 00000000 00000000 | ..... |
:
:
SPTAE: 0295721C
  SPTFLAGS. 10          SPTALLOC. 00000000 SPTFREE.. 02C53008
  SPTSIDEQ. 00000000 SPTUSECT. 00000320 SPTHIUSE. 000003C0
  SPTNBRPG. 00000001 SPTLNPTH. 000000A0
:
:
  DATA: 0295721C
    +0000 02957204 0295725C 00000000 00000000 | .n...n.*..... |
    +0010 00000000 00100000 00000000 02C53008 | .....E.. |
    +0020 00000000 00000320 000003C0 00000001 | .....{.... |
    +0030 00000000 000000A0 00000199 00000001 | .....μ...r... |
:
:
SPTAE: 0295725C
  SPTFLAGS. 00          SPTALLOC. 00000000 SPTFREE.. 00000000
  SPTSIDEQ. 00000000 SPTUSECT. 00000000 SPTHIUSE. 00000000
  SPTNBRPG. 00000000 SPTLNPTH. 00000178
:
:
  DATA: 0295725C
    +0000 02957204 00000000 00000000 00000000 | .n..... |
    +0010 00000000 00000000 00000000 00000000 | ..... |
    +0020 00000000 00000000 00000000 00000000 | ..... |
    +0030 00000000 00000178 0000000A 00000000 | ..... |

```

TOPOLOGY

Use TOPOLOGY to help diagnose topology and routing problems. The TOPOLOGY provides the summary information output of user-selected criteria for the control blocks representing node records and TG records.

Operands

COSNAME

The name of the entry in the APPN Class of Service table should be 1-8 alphanumeric characters. If a *cosname* value is not specified, the default class of service name is #CONNECT.

ALLTOPO

Displays a summary of all node records and a summary of all TGs that originate at each node. Displays the weights of all node records and TG records.

This is the default value.

ANNTOPO

Displays a summary of all network node records and a summary of all TGs that originate at each network node. Displays the weights of all network node records and TG records.

AENTOPO


Displays a summary of all end node records and a summary of all TGs that originate at each end node. Displays the weights of all end node records and TG records.

ACNTOPO

Displays a summary of all connection network node records and a summary of all TGs that originate at each connection network node. Displays the weights of all connection network node records and TG records.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

TOPOLOGY ALLTOPO

ALL Nodes Topology Summary

NodeName	NDRECAAdr	Type	SeqNo	GCI	QUIES	HPR	Weight		
NETSOUTH.GERMANY	3EC2B4D8	EN	00000000	N	N	NO	160		
Total Number of TG Records Found:			0						
NETSOUTH.RUSSIA	3EC2B670	NN	00000000	N	N	NO	60		
Total Number of TG Records Found:			0						
NETMAP3.LVRN6E	3EC2BCD0	CN	00000000	N	N	NO	0		
Total Number of TG Records Found:			0						
NETSOUTH.SPAIN	3EC2B340	EN	00000000	N	N	NO	160		
Total Number of TG Records Found:			0						
NETSOUTH.FRANCE	3EC2B9A0	EN	00000000	N	N	NO	160		
Total Number of TG Records Found:			0						
NETMAP3.GVRN4A	3EFF81A8	CN	00000000	N	N	NO	0		
Total Number of TG Records Found:			0						
NETSOUTH.BOTSWANA	3EC2B010	MDH	00000006	N	N	CONTR	60		
Dest CPNAME	TGRECAAdr	TGN	STAT	SeqNo	GCI	QUI	HPR/TT	TGTYPE	Weight
NETSOUTH.SPAIN	3F031010	21	INOP	00000002	N	N	Y/Y	ENDPT	32767
NETSOUTH.GERMANY	3F0312D0	21	INOP	00000002	N	N	Y/Y	ENDPT	32767
NETSOUTH.RUSSIA	3F031590	21	INOP	00000006	N	N	Y/Y	ENDPT	32767

TRSTRACE

```
+0090 C1000000 000000D5 C5E3C14B E2E2C3D7 | A.....NETA.SSCP |
+00A0 F4C30000 00000000 E3C7E400 00000000 | 4C.....TGU..... |
      |           |           |           |           |           |
+0F90 C1000000 000000D5 C5E3C14B E2E2C3D7 | A.....NETA.SSCP |
+0FA0 F5C10000 00000000 E3C7E400 00000000 | 5A.....TGU..... |
+0FB0 0000000A BD61BF07 BD61BF1D 16B0D850 | ...../.../....Q& |
+0FC0 D5E30710 2120D5C5 E3C14BE2 E2C3D7F1 | NT....NETA.SSCP1 |
+0FD0 C1000000 000000D5 C5E3C14B E2E2C3D7 | A.....NETA.SSCP |
+0FE0 F5C10000 00000000 E3C7E400 00000000 | 5A.....TGU..... |
+0FF0 0000000A BD61BF07 BD61BF1D 16B122D0 | ...../.../.....} |
TRACPAGE: 16A7B020
+0000 D5E30810 2120D5C5 E3C14BE2 E2C3D7F1 | NT....NETA.SSCP1 |
+0010 C1000000 000000D5 C5E3C14B E2E2C3D7 | A.....NETA.SSCP |
+0020 F5C10000 00000000 E3C7E400 00000000 | 5A.....TGU..... |
```

VITAL

Use the VITAL function to extract an internal VIT from a dump for use with the VIT analysis tool. See Chapter 8, "Using the VIT analysis tool," on page 365.

Before using VITAL, allocate a data set for the specified DD name. The data set, when VITAL is invoked, must have these attributes.

```
RECFM=VB
LRECL=284
DSORG=PS
```

Note:

1. The VIT extracted by VITAL can be used as input to the VIT analysis tool, but not to the IPCS GTFTRACE subcommand.
2. Time stamps are provided with each entry in the VITAL output, but it should be understood that these are approximated time stamps generated during the extraction process.
3. It is possible for the extracted VIT to contain one or more entries that begin with one or more words of hexadecimal zeros.
 - While VTAM does not create such VIT entries, it is possible for user programs to create these non-standard VIT entries. These entries are processed as is by the VITAL function.
 - It is possible for the end of the VIT table to contain some trace entries which are all zeros. These entries are extracted by the VITAL function as non-standard trace entries.

Operands

DD name

Specify the name of the DD statement allocated to receive the extracted VIT. The DD name should be a 1–8 alphanumeric character name. If it contains fewer than 8 characters, it is padded on the right with blanks.

The DD name is required.

You must allocate the specified data set before VITAL is invoked. VITAL will not allocate the data set for you.

Note: The jobname field in the GTF header is set to VFDTRACE. The ASCB address field is set to 0.

Use the following syntax as an alternative to the panel interface.

Syntax

```
▶▶—VITAL— —DDN—(—DD_name—)————▶▶
```

Sample output

VITAL

VITAL DD(VITDATA)

VITAL Analysis

VITAL processing completed successfully

VTAM

Use VTAM to format and display the following information:

- RDT and RDTEs
- Memory process scheduling table (MPST) and process scheduling table (PST)
- ACDEBs, APPCBs, COPRs, FMCBs, FMCBEXTs, HSICBs, LUCBs, NSICBs, NSSCBs, RABs, and SABs
- NCBs
- Buffer pool control blocks (BPCBs), buffer pool directory (BPDTY), pool extension blocks (PXBs), storage pool anchor block (SPANC), SPANC task-associated element (SPTAE), storage pool page table (PAGTB), and storage pool page table entries (PTEs)
- Locked queue anchor block (LQAB)
- Waiting request elements (WREs) and event identifiers (EIDs)
- Modules from the ATCVT, in the form *module name* and *module address*, sorted by module name

Use the following syntax as an alternative to the panel interface.

Syntax

```
▶▶—VTAM————▶▶
```

Sample output

VTAM

VTAM Analysis

BPD

DATA: 02953000

+0000	000C000C	00000000	02953508	00000000	n.....	
+0010	00000000	00002000	7FFFFFFF	02953390	".....n..	
+0020	029531A8	00CC41F8	00000000	000003E8		.n.y." .8.....Y	

⋮

Buffer pool ID SMS1

BPCB: 02953390

BPCBRPHA.	00000000	BPCBRPHB.	00000000	BPCBRPH1.	00000000
BPCBRPH2.	00000000	BPCBAVNO.	00000000		

DATA: 02953390

+0000	00000000	00000000	600000E7	02954FF8	-..X.n 8	
+0010	00000000	00000000	00000000	00000000		

VTAM

```

+0020 00000000 00000000 02953000 00000000 | .....n..... |
:
SPANC 02958204
POOLNUM 0000      POOLNAME RUPEPRIV ASSOCID N/A
DATA: 02958204
+0000 D9E4D7C5 D7D9C9E5 000C0002 00000000 | RUPEPRIV..... |
+0010 00000000 00000000 | ..... |
SPTAE: 0295821C
SPTFLAGS. 10      SPTALLOC. 00000000 SPTFREE.. 02C4F008
SPTSIDEQ. 00000000 SPTUSECT. 000000A0 SPTHIUSE. 00000500
SPTNBRPG. 00000001 SPTLNTH. 000000A0
DATA: 0295821C
+0000 02958204 0295825C 00000000 00000000 | .nb..nb*..... |
+0010 00000000 00100000 00000000 02C4F008 | .....D0. |
+0020 00000000 000000A0 00000500 00000001 | .....μ..... |
:
MPST: 00CBED38
MPSCHAIN. 00CBEDB8 MPSPSTQ.. 02818328
DATA: 00CBED38
+0000 D4D7E2E3 80000000 00CBEDB8 00700016 | MPST.....½.... |
+0010 00F74180 02818328 0290EB18 00AFB040 | .7...ac....._[ |
+0020 C1E4E3C8 823BA7CA 00FDD6F6 00160002 | AUTHb.x...06.... |
:
PST: 02818328
+0000 61000480 00000000 02818328 00CBED38 | /.....ac..... |
+0010 00000000 00000000 00000000 00000000 | ..... |
+0020 00000000 00000051 00AF8FF0 00000000 | .....±0.... |
:
QAB: 02955740
+0000 D9C4E340 000D0000 00CC1AB0 02CE0008 | RDT .....".[".. |
+0010 00100000 00700074 | ..... |
RDT: 00CC1AB0
RPRNAME.. VTAMSEG RPRENTRY. 02      RPRBITAN. 02000100 00
RPRDEVCH. 00000000 00000000
DATA: 00CC1AB0
+0000 E5E3C1D4 E2C5C740 80000000 00020200 | VTAMSEG ..... |
+0010 00000000 00000000 00000000 00CC1B18 | .....". |
+0020 00000000 000000B0 00000000 00000000 | .....[..... |
:
RDTE: 00CC1B60
RPRNAME.. SSCP1A RPRENTRY. 11      RPRBITAN. 00000940 00
RPRDEVCH. C06D0000 00800000
DATA: 00CC1B60
:
ATCVT: 00CC41F8
ISTACC00. 82A957A0 ISTACC01. 82ACD8B8 ISTAICIR. 80DD9000
ISTAICPT. 823B7014 ISTAPCAD. 823B8298 ISTAPCES. 823BC560
ISTAPCGT. 823C5078 ISTAPCIE. 822E29F0 ISTAPCIN. 823BD990
ISTAPCKU. 80DDA5A0 ISTAPCPC. 823BCE58 ISTAPCPD. 823B8920
ISTAPCPS. 823C0580 ISTAPCRP. 822E2934 ISTAPCRS. 823B8F84
:
ATCIOLQB
LQAB: 02C6DF30
LQABFRST. 00000000 LQABLAST. 00000000 LQGSUBA.. 00000001
DATA: 02C6DF30
+0000 00000000 00000000 02D3D8C7 00000001 | .....LQG.... |
+0010 00000000 00000000 | ..... |
ATCLUSMQ
LQAB: 02A3BC14
LQABFRST. 00000000 LQABLAST. 00000000
DATA: 02A3BC14
+0000 00000000 00000000 03D3D8C7 | .....LQG |
ATCMCQAB
LQAB: 02A3BC08
LQABFRST. 00000000 LQABLAST. 00000000
DATA: 02A3BC08

```

```

+0000 00000000 00000000 01D3D8C7      | .....LQG |
ATCPULQB
LQAB: 02A3BCB4
LQABFRST. 00000000 LQABLAST. 00000000
DATA: 02A3BCB4
+0000 00000000 00000000 04D3D8C7      | .....LQG |
NODAT_CPWAIT_QUEUE
LQAB: 02A3BCC0
LQABFRST. 00000000 LQABLAST. 00000000
DATA: 02A3BCC0
+0000 00000000 00000000 05D3D8C7      | .....LQG |
ATCSSLQB
LQAB: 02A3BCCC
LQABFRST. 00000000 LQABLAST. 00000000
DATA: 02A3BCCC
+0000 00000000 00000000 06D3D8C7      | .....LQG |
ATCSSMQB
LQAB: 02A3BD6C
LQABFRST. 00000000 LQABLAST. 00000000
DATA: 02A3BD6C
+0000 00000000 00000000 07D3D8C7      | .....LQG |

```

VTBASIC

Use the VTBASIC function to display the ATCVT, the configuration table (CONFT), the component recovery areas (CRAs), and the VTAM internal trace (VIT).


Operands

Trace output

Enter **Format** to format the VIT and **No format** to display the VIT in hexadecimal format. **Format** is the default.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

VTBASIC NOFORMAT

```

VTBASIC Analysis
VTAM INTERNAL TRACE TABLE 000001EF_81000000

PRESENT WRAP C728E7D3 D61C13C4 LAST WRAP 00000000 00000000
CURRENT ENTRY 000001EF_83503020 LAST ENTRY 000001EF_841FFFE0
C4E2D740 12582410 02915E88 00CC4908 02A275F8 02A275F8 E3E2E6E4 02929010
D3D2E2C8 12000100 00CC4C70 00000000 82A95442 00000000 00000000 02929010
E4D5D3D2 12000100 00CC4C70 00000100 82A9546C 00000000 01000000 02929010
D8E4C558 12482810 02915E88 00CC4248 82A954F8 02A275F8 C9D5E3D4 02929010
C5E7C9E3 12000010 02915E88 00CC4908 82A95668 80000000 E3E2E6E4 02929010
D9C5D8E2 12170000 02915E88 02928810 822E2E16 00010000 0290DDA0 00000000
C4E2D740 12582810 02915E88 00CC4248 02A275F8 02A275F8 C9D5E3D4 02928810
C5E7C9E3 12000010 02915E88 00CC4248 82A22B1C 80000000 C9D5E3D4 029 28810
D9C5D3E2 12170000 02915E88 02928810 822E2CB6 00000000 0290DDA0 00000000

```

VTBASIC

```
⋮
ATCVT: 00CC41F8
  ATCRDT... 02955740  ATCSRT... 02C35008  ATCCONFT. 00CC18E8
  ATCBPDA.. 02953000  ATCACTRM. 0000      ATCVTL0D. 02A27650
  GWSSCP = YES
DATA: 00CC41F8
+0000 E5C5F4F3  40404040  FFF900C8  02825000  | VE43   .9.H.b&  |
+0010 00000000  0000FFF9  11280000  00000000  | .....9.....  |
+0020 02915E88  00000000  00000000  00000000  | .j;h.....  |
+0030 00CC4524  00000000  13201000  00000010  | .".....  |
+0040 11280000  00000000  02915E88  00000000  | .....j;h....  |
⋮
+07B0 00000000  00000000  00000000  00000000  | .....  |
+07C0 00000000  00000000  00000000  00000000  | .....  |
+07D0 00000000  00000000  00000000  00000000  | .....  |
+07E0 00000000  00000000  00000000  00000000  | .....  |
```

VTBUF

Use VTBUF to analyze buffer pool control blocks (BPCBs) and obtain a status summary for all buffer pools or a specific buffer pool. For each buffer pool, the following information is displayed:

- Starting and ending address of buffer pools
- Buffer pool address (BPCB)
- Buffer type (fixed or pageable)
- Buffer size
- Number of buffers allocated
- Slowdown threshold
- Number of buffers available
- Expansion threshold
- Contraction threshold
- Number of times expanded
- Maximum number of buffers
- Expansion increment
- Expansion size
- Percentage of buffers in use
- Total number of buffers
- Bytes in static and expanded areas
- Buffers in other pools
- Buffers in static area
- Total queued request parameter headers (RPHs)

If any expansions have occurred, the pool extension block (PXB) address, number of buffers available, totals buffers, beginning of the extent, and the first available extent are also presented.

Operands

Buffer name

Enter a 2-7 character buffer name in the form *cc* or *cccc*BUF where *cccc* is the buffer name. Counts and totals information will be displayed in decimal form.

The default is ALL.

Equated symbol

Symbol

Description

buffername **BSTART**

Each starting buffer address


buffername **BEND**

Each ending buffer address

Note: For information on the DISPLAY BFRUSE (buffer use) command, which displays information about VTAM buffer use, see z/OS Communications Server: SNA Operation.

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

VTBUF BUFFER(IO)

```

                                VTBUF Analysis
                                IO Buffer Analysis
Size of Buffers(bytes)          345  Buffer maximum                110
Total buffers available         94   Static buffers allocated    110
Total number of buffers        110
Buffers in use (%)              15   Available static buffers    94
Bytes in static &
  expanded areas                37950
Slowdown threshold              5   Expansion threshold         6
Contraction threshold          32767
Number of expansions            0   Expansion increment         11
Expansion size                  4096
Total queued RPHs              0
Fixed or pageable?             FIXED
Buffer pool address             X'028DB410'
Beginning address of pool       X'028F9000'
Ending address of pool          X'02903000'
Buffer pool has no extensions
  
```

VTCVTPAB

Use VTCVTPAB to obtain a list of the PABs and DYPABs in the ATCVT. The PAB control block for each PAB is formatted. For the very extended PABs, the address of the first element on the PAB is displayed. The following PABs and DYPABs are processed:

- ATCCSPAB - Configuration services DYPAB
- ATCITPAB - Internal trace DYPAB
- ATCLUSRT - LU services router PAB
- ATCNSPAB - TSC no session PAB
- ATCPDPAB - Problem determination trace PAB

VTCVTPAB

- ATCPUIOP - SSCP/PU services I/O PAB
- ATCPUPAB - PU services PAB
- ATCPXPAB - Dynamic buffer pool expansion PAB
- ATCRYPAB - Definition for CRYPTO task
- ATCSOPAB - Session outage notify PAB
- ATCSSPAB - Session serialization PAB
- ATCTMRPB - Timer services DYPAB
- ATCTPMPB - CPMSG PAB
- ATCTRMPB - Termination task
- ATCVDPAB - SSCP VARY definition PAB
- ATCWUPAB - Wake up PAB
- NODAT_NOSPAB - Network operator services PAB

Use the following syntax as an alternative to the panel interface.

Syntax

▶—VTCVTPAB—▶

Sample output

VTCVTPAB

VTCVTPAB Analysis

```
Configuration Services elements
PAB: 00C17278
+0000 86385C68 80000000 065E4100 00000000 | f.*.....;..... |
+0010 10321000 00000010 00000000 00000000 | ..... |
+0020 00000000 00000000 00000000 00000000 | ..... |
      1 level elements          0
      2 level elements          0
      3 level elements          0
      4 level elements          0
      5 level elements          4    07524DE0
      6 level elements          55   073BAA60
      7 level elements          45   0751FAC0
Definition for Crypto Task elements          0
PAB: 00C17CC8
+0000 00000000 00000000 00C17CA4 00000000 | .....A@u.... |
+0010 60001000 00000010 | ..... |
Dynamic Buffer Pool Expansion elements          0
PAB: 00C172E8
+0000 00000000 00000000 00C17D84 06312010 | .....A'd.... |
+0010 0E0C1000 00000010 | ..... |
VTAM Termination Task elements          0
PAB: 00C17D30
+0000 00000000 00000000 00C17DAC 00000000 | .....A'..... |
+0010 0D201000 00000010 | ..... |
Internal Trace elements          0
PAB: 00C17220
+0000 00000000 00000000 00C17524 00000000 | .....A..... |
+0010 13201000 00000010 | ..... |
LU Services Router elements
PAB: 00C17380
```



```

+0000 86385D40 80000000 00C174DC 06312810 | f.) .....A..... |
+0010 0F321000 00000010 00000000 00000000 | ..... |
+0020 00000000 00000000 00000000 00000000 | ..... |
  1 level elements          0
  2 level elements         0
  3 level elements         0
  4 level elements         1    06204708
  5 level elements         9    06204CA8
  6 level elements         8    073BA920
  7 level elements         0

Network Operator Services elements          0
PAB: 00C17408
+0000 00000000 00000000 00C174B4 00000000 | .....A..... |
+0010 14201000 00000010 | ..... |

PD Trace elements                          0

PAB: 00C17B18
+0000 00000000 00000000 00C17B04 00000000 | .....A#..... |
+0010 12201000 00000010 | ..... |

PU Services elements
PAB: 00C17310
+0000 06385E18 00000000 06572F38 00000000 | ..;..... |
+0010 06321000 00000010 00000000 00000000 | ..... |
+0020 00000000 00000000 00000000 00000000 | ..... |
  1 level elements          0
  2 level elements         0

Session Outage Notify elements              0
PAB: 00C17458
+0000 00000000 00000000 00C17650 00000000 | .....A.&.... |
+0010 28201000 00000158 | ..... |

Session Serialization elements              0
PAB: 00C17440
+0000 00000000 00000000 00C1764C 00000000 | .....A.<.... |
+0010 27001000 00000140 | ..... |

VARY Definition elements                   0
PAB: 00C172B8
+0000 00000000 00000000 065A18D4 00000000 | .....!.M.... |
+0010 0B211000 00000010 00000000 00000000 | ..... |

SSCP/PU Services I/O elements              0
PAB: 00C17350
+0000 80000000 80000000 00C17750 06316010 | .....A.&...- |
+0010 11211000 00000010 00000000 00000000 | ..... |

Timer Services elements                    0
PAB: 00C17248
+0000 00000000 00000000 00C174CC 00000000 | .....A..... |
+0010 0C281000 00000010 | ..... |

TPMSG elements                             0
PAB: 00C179C8
+0000 80000000 80000000 00C175A8 06383810 | .....A.y.... |
+0010 17251000 00000010 00000000 00000000 | ..... |

TSC No Session elements                    0
PAB: 00C173C0
+0000 00000000 00000000 00C17548 00000000 | .....A..... |
+0010 26001000 00000010 | ..... |

Wakeup PAB elements                        0

```

```
PAB: 00C17908
      +0000 00000000 00000000 00C1792C 06313010 | .....A^..... |
      +0010 29241000 00000010 | ..... |
```

VTFNDMOD

Use VTFNDMOD to determine the VTAM module name and displacement for a given address.

VTFNDMOD is useful for converting the issuer address (ISSR) in a VIT into a module name and displacement. It searches up to 5000 bytes before the specified address.

In addition to the module name and displacement into the module that the specified address represents, the following information is displayed:

- Date module compiled
- PTF level, if any
- Address entered
- Module entry point address
- Address displacement into module
- First hexadecimal 40 bytes of the module
- Hexadecimal 40 bytes around the entered address
- The beginning and ending address of a region, if the address is in FLPA, MLPA, PLPA, extended PLPA, extended FLPA, or extended MLPA

Operands

You must specify one address or one symbol.

Address

Enter 1–8 hexadecimal digits in the form X'x...' for the address used to determine the VTAM module name and displacement. If the address is fewer than 8 digits, it is padded on the left with zeros.

Note: The address must be located after the module name at the start of the module.

IPCS symbol

Enter 1–31 alphanumeric characters for an IPCS symbol name that has been previously equated to a location within a VTAM module. Do not include a period.

Under IPCS, the symbol X represents the address currently being displayed. If the current address points to a location within a VTAM module, X may be used to refer to it.

Equated symbol

After the module name is determined, an IPCS symbol (the module name) is equated to the beginning of the CSECT.

Symbol

Description

module eye-catcher

Module entry point

Use the following syntax as an alternative to the panel interface.

Syntax

```

▶ VTFNDMOD [ ADDR (address) | SYMBOL (symbol) ]

```

Sample output

```
VTFNDMOD ADDR(X'2B023C0')
```

VTFNDMOD Analysis

```
Module name:          ISTOREI
Compile date:        92.262
```

```
Address entered:     02B023C0
Module entry point:  02B022C0
                   -----
```

```
Displacement into module: 100
```

First '40'X bytes of module:

```

DATA: 02B022C0
+0000 47F0F014 0FC9E2E3 D6D9C3C5 C940F9F2 | .00..ISTORCEI 92 |
+0010 4BF2F6F2 90ECD00C 18CF41B0 CFFF41A0 | .262..}..... |
+0020 BFFF4190 AFFF4180 9FFF4170 8FFF50D0 | .....&} |
+0030 C5C84160 C5C45060 D00818D6 1F005000 | EH.-ED&-};.0.&. |

```

Storage around address entered:

```

DATA: 02B023AC
+0000 C54A5860 04085800 65701860 D203602C | E+.-.....-K.-. |
+0010 C61CD203 6028C618 5800C630 18204100 | F.K.-.F...F.... |
+0020 00585830 20005030 60001E06 8B300002 | .....&.-..... |
+0030 1E035000 C7085820 C7145020 60485020 | ..&;G...G.&.-.&. |

```

VTMODS

Use VTMODS to find the entry point of the VTAM modules that reside in the VTAM private region. VTMODS reports the number of modules found and equates the entry point of each module found to its module name in the IPCS symbol table. After VTMODS executes, the VTAM modules found can be located using the module name in the IPCS LIST command.

VTMODS is useful when you are checking the PTF level of several modules or when you want to quickly verify the PTF or APAR level of a module in a dump.

Note: For VTMODS to execute successfully, the VTAM private region must have been dumped (that is, the RGN parameter must have been specified when the dump was taken). If fewer than 10 modules are found, VTAM private storage is missing from the dump.

If the dump does contain the VTAM private region but is a partial dump, VTMODS attempts to find as many VTAM modules as possible.

Operands

List

The default is N. Specify Y to receive a list of each module found, its entry

point address, compile date, and PTF level, if present. The modules are in the order that they were found in storage (that is, by storage address, lowest to highest), followed by a list of the modules in alphanumeric order.

Equated symbol

For each VTAM module that is found, an IPCS symbol (the name of the module as it appears in the module eye-catcher) is equated to the entry point of the module.

Symbol

Description

module eye-catcher

Module entry point

Use the following syntax as an alternative to the panel interface.

Syntax



Additional information

The symbols created remain in the IPCS dump directory until the dump directory is deleted or until an IPCS DROPDUMP command is issued for the dump.

VTAMODS scans VTAM private storage for the character string IST. When IST is found, a check is made to determine whether this occurrence of IST represents a VTAM module eye-catcher. Most VTAM modules have a branch instruction hexadecimal 47F0F0xx at the entry point to the module to branch around the eye-catcher. If hexadecimal 47F0F0xx appears 5, 7, or 9 bytes before the module eye-catcher, the location is considered in most cases to be a VTAM module entry point.

Storage is scanned starting at the lowest address of private storage to the top of the private region (below the 16 MB line). The scan then continues starting at the lowest address of extended private storage (above the 16 MB line) and continues for about hexadecimal 400 000 bytes.

Sample output

VTAMODS LIST(Y)

Address	Module	VTAMODS Analysis	
		Compiled	PTF
-----	-----	-----	-----
000063CA	ISTATM00	91.322	
00006A70	ISTINCBX	91.322	
00006CD0	ISTINCRS	91.320	
000072E0	ISTIECHS	91.320	
0000772A	ISTCPM01	91.319	
00007768	ISTSSCX	91.329	
00009DC8	ISTCPCIT	91.322	
0000A488	ISTINCR4	91.322	
0000B7E8	ISTPUCWI	91.320	
0000BD10	ISTSSCX	91.320	
0000C0F0	ISTCSCEX	91.322	
0000CDA0	ISTPDCLU	91.322	

0000EAA8	ISTPDCSE	91.320
00011588	ISTPUCX0	91.320
000124B0	ISTCSCSD	91.319
00015C00	ISTENQIO	91.336
00015EC0	ISTENQPR	91.326
00016A18	ISTENQIN	91.320
00016D90	ISTENQRT	91.320
02A00138	ISTINM01	91.338
02A01C68	ISTLUCQD	91.320
02A01EB0	ISTCPCQD	91.353
02A021B0	ISTPUCQD	91.319
02A021F8	ISTINCXR	92.003
02A02788	ISTSCCIT	92.002
02A02BC8	ISTINCCT	91.352
02A02EA8	ISTINCIT	91.352
02A03100	ISTINCPD	91.352
02A03E78	ISTINFIC	91.352
02A04680	ISTCICPR	91.352
02A082A0	ISTCICDF	91.352
02A08968	ISTINCCP	91.350
02A094E0	ISTORCEI	91.346
02A0EE88	ISTCICTR	91.346
02A10DA0	ISTINCTR	91.346
02A11EC0	ISTINCCF	91.346
02A132F0	ISTDRCIT	91.344
02A13B50	ISTINCSA	91.344
:		
02C31168	ISTXP1WB	91.351
02C31848	ISTXP1WC	91.351
02C31D48	ISTXP1WR	91.351
02C32748	ISTXP1WS	91.351

VTNODE

Use VTNODE to determine:

- If an SIB exists on the secondary chain where the RDTE is the SLU
- If an RDTE application exists, and if the ACDEB, LUCB, FMCB, and FMCB extension associated with the session exist

If any SIBs exist on the secondary chain, only the first SIB is processed. To process the PLU, use the SES function. Excerpts of the SIB, LU RDTE, APPL RDTE, ACDEB, LUCB, FMCB, and FMCB extension are displayed if present.

Operands

RDTE name

The RDTE name of a CDRSC or LU RDTE should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks. The entered RDTE name must be the SLU, and the PLU must be an application. For CDRSC independent LUs, VTNODE will process only the first entry.

The RDTE name is required.

NetID

The NetID name representing the network ID of another network outside the host network where a resource resides should be 1–8 alphanumeric characters. If it contains fewer than 8 characters, it is padded on the right with blanks.

Use the following syntax as an alternative to the panel interface.

Syntax

```

▶—VTNODE— —RDTE—(—name—)—————▶
                                     |
                                     | —NETID—(—netid—)

```

Sample output

VTNODE RDTE(TCPM1010)

```

VTNODE Analysis
RDTE: 09F90590
RPRNAME.. TCPM1010 RPRENTRY. 83 RPRBITAN. 01000110 00 RPRDEVCH. C06D0000 00800000
DATA: 09F90590
+0000 E3C3D7D4 F1F0F1F0 80000000 00830700 | TCPM1010.....c.. |
+0010 00000000 00000000 00200000 09F90640 | .....9. |
+0020 09F905F8 FFFFEC0 FFFFEC0 00000000 | .9.8...{...{.... |
:
SIB: 0A061300
SIBFSMIN. FC SIBFSMTM. 00 SIBFSENS. 00000000 SIBBPRIQ. 00000000 SIBBSEQ. 00000000 SIBTMFL. 00
SIBTREAS. 00 SIBTSESE. 00
DATA: 0A061300
+0000 9800FC00 00000000 EAABEEC3 939BA822 | q.....C1.y. |
+0010 40404040 00000000 40404040 D5E2E7F3 F2F7F0F2 | .....NSX32702 |
+0020 D5E2E7F3 F2F7F0F2 B33C21C8 C81A1B07 | NSX32702...HH... |
:
RDTE: 09F97BD0
RPRNAME.. APPC2A02 RPRENTRY. 55 RPRBITAN. 01000910 01 RPRDEVCH. C06D0000 00800000
DATA: 09F97BD0
+0000 C1D7D7C3 F2C1F0F2 80000000 00550200 | APPC2A02..... |
+0010 40040000 00020049 00200000 09F97C80 | .....9e. |
+0020 09F97C38 00000120 00000120 00000000 | .9e..... |
:
ACDEB: 08C8C6A8
ACDTCB... 006E5A70 ACDCHN... 00000000 ACDRDTE.. 09F97BD0
DATA: 08C8C6A8
+0000 0F480000 00000000 08B587C8 00000000 | .....gH... |
+0010 00000000 00000000 009F83A0 00000000 | .....c..... |
+0020 2D010000 0FF00010 00000000 00000000 | .....0..... |
:
LUCB: 08D24110
+0000 52700049 00000000 08B587C8 00000000 | .....gH... |
+0010 00000000 80000000 08B47018 01FD0000 | ..... |
+0020 00040000 00000000 00000000 09F97BD0 | .....9#} |
:
No TSCB elements queued to LUCB PAB
FMCBX: 08C63378
+0000 00000002 001F8402 01000003 08B2D018 | .....d.....}. |
+0010 00000002 00910000 0002001F 08C63258 | .....j.....F.. |
+0020 08C63408 08C645D8 08C648B8 00000000 | .F...F.Q.F..... |
:
FMCB: 08B2D018
+0000 03408002 08C63378 08B587C8 00000000 | . ...F...gH... |
+0010 08C6C210 00000000 00000000 00000000 | .FB..... |
+0020 08C6B260 00000000 1C016200 0FF00018 | .F.....0.. |
:

```

VTREADYQ

Use VTREADYQ to analyze some of the major control blocks associated with an application. For each memory process schedule table (MPST) chain, the VTAM data extent control blocks (ACDEBs), process scheduling tables (PSTs), logical unit control blocks (LUCBs), and function management control blocks (FMCBs) are checked for PABs, DYPABs, and ready queues that contain queued elements. VTREADYQ lists the first elements on those queues and PABs.

Note: For a large network, this could take several minutes to run.

Use the following syntax as an alternative to the panel interface.

Syntax

▶▶—VTREADYQ—◀◀

Sample output

VTREADYQ

```

                                VTREADYQ Analysis
MPST      1
Processing begins on FMCB extension for LU TSO0001
MPST      2
Processing begins on FMCB extension for LU TSO
MPST      3
Processing begins on FMCB extension for LU APPCAP06
Processing begins on FMCB extension for LU APPCAP05
MPST      4
Processing begins on FMCB extension for LU ISTATAO0
There are no FMCB extensions off of LUCB 02924090
Processing begins on FMCB extension for LU ISTPDCLU
Synchronous TPROSTed RPH count          1
ELEMENT: 8295B500
+0000  01C40020  80000000  02915E88  00CC45AC  |.D.....j;h.".-|
+0010  82A332A4  6C000010  00CC49B8  00000000  |bt.u%. ....".½....|
+0020  00000000  0295B500  00000000  0295B500  |.....n.....n_.|
+0030  80000000  00000000  00CC41F8  00000041  |.....".8....|
+0040  00000000  00000300  00000000  80000000  |.....|
+0050  00000000  00000000  82A32800  02A334B0  |.....bt...t.[|
+0060  00CC41F8  00000000  00000000  00000000  |."8.....|
+0070  00000000  00000000  00000000  00000000  |.....|
+0080  00000000  00000000  00000000  00000000  |.....|
+0090  00000000  00000000  00000000  00000000  |.....|
+00A0  00000000  00000000  00000000  00000000  |.....|
+00B0  00000000  00000000  00000000  00000000  |.....|
+00C0  00000000  |.....|
Synchronous normal PAB          1
ELEMENT: 00CC4248
+0000  80000000  80000000  00CC44CC  00000000  |.....".". ....|
+0010  0CA81000  00000010  |.y.....|
Processing begins on FMCB extension for LU SSCP1A
Processing begins on FMCB extension for LU VTAM

```

VTRPH

Use VTRPH to analyze the entire LP buffer pool of request parameter headers (RPHs) and display those that are waiting, running, holding locks, or are in error.

If an RPH is waiting at an address other than X'0' or X'FFFFFFFF', the major control block and the current process anchor block (PAB) are listed. In addition, the resume addresses are shown with the number of RPHs that were waiting at those addresses.

Use the following syntax as an alternative to the panel interface.

Syntax

▶▶—VTRPH—◀◀

Sample output

VTRPH

```

                                VTRPH Analysis
                                LP Buffer Analysis
Buffers available                56
Total number of buffers        64
Number of expansions            0
Buffer found does not contain an RPH at address X'02928010'
RPH at buffer address X'02929010' is running, RPHRESUM = 0
Module was not found
RPH Major control block:
  DATA: 00CC48F8
          +0000 11280000 00000000 02915E88 00000000 | .....j;h.... |
          +0010 00000000 00000000 00CC492C 02929010 | .....".k.. |
          +0020 29241000 00000010 | ..... |
Work elements found            4
RPH work element address X'02A275F8'
RPH at buffer address X'0292B010' is running, RPHRESUM = 0
Module was not found
RPH Major control block:
  DATA: 00CC43F8
          +0000 11280000 00000000 02915E88 00000000 | .....j;h.... |
          +0010 80000000 80000000 00CC44B4 0292B010 | .....".+.k[. |
          +0020 14201000 00000010 | ..... |
Work elements found            1
RPH work element address X'02C4F340'
Error buffers found            1
Unallocated buffers found      61
Total number of buffers processed 64
No allocated CRAs were found

```

VTVIT

Use VTVIT to determine which VIT options were in effect at the time of a dump, and whether the trace was running internally (MODE=INT), externally (MODE=EXT), or internally and externally.

If the VIT was running externally, no further processing occurs. IPCS symbols are created for the beginning and end of the internal VIT table, and for the current, oldest, and last VIT entries.

An option is available to produce an unformatted listing of the entire VIT table. Use VTBASIC to format the VIT table.

To extract a VIT from a dump for use with the VIT analysis tool, use VITAL.

Operands

Search argument

Scan displays the VIT entries containing a specified search argument. Enter 1–8 alphanumeric characters or 1–16 hexadecimal digits in the form X'x...' for the search argument.

If the hexadecimal data string is not an even number of digits, the high-order half-byte is set to 0.

List VIT

The default is N. Specify Y to list the entire VIT table. The internal VIT table is processed in the following order:

1. From the oldest trace entry in the trace table in 64-bit common (HVCOMMON) storage to the end of the trace table.
2. From the beginning of the trace table in 64-bit common (HVCOMMON) storage to the current entry.


This results in “unwrapping” the trace table so the trace entries are processed and listed in chronological order (that is, the oldest trace entry is listed first at the top of the output, and the newest trace entry is listed last at the bottom of the output). A message is included in the output to indicate where the physical end of the trace table was encountered.

Equated symbol

Symbol	Description
VIT	The beginning of the VIT table
VITC	The current VIT entry
VITO	The oldest VIT entry
VITE	The end of the VIT table.

Use the following syntax as an alternative to the panel interface.

Syntax



Additional information

The beginning of the trace table is mapped by ITTRC in z/OS Communications Server: SNA Data Areas Volume 1. It contains the present-wrap time stamp, last-wrap time stamp, current-entry address, and last-entry address. The current entry is the most recent entry (that is, the last entry to be written before the dump was taken). The last entry is the one that was written in the last position of the in-storage trace table before wrapping to the beginning of the trace table.

Sample output

```

VTVIT

                                VTVIT Analysis

VTAM external trace options active at the time of this dump

API LOCK PSS SMS PIU MSG SSCP CIO NRM APPC VCNS LCS

VTAM internal trace options active at the time of this dump

API PIU MSG SSCP NRM

Pages in VTAM internal trace table (Decimal) =                12800

VIT - Start of VTAM internal trace table:      000001EF_81000000
VITC - Current VTAM internal trace table entry: 000001EF_83503020
VITO - Oldest VTAM internal trace table entry: 000001EF_81000040
VITE - End of VTAM internal trace table:      000001EF_84200000
VTAM internal trace table (oldest to newest entry)

```

VTVIT

```
VITPAGE: 000001EF 81000040
+0000 D8E4C500 1F480110 06337250 063372A0 | QUE.....&....
+0010 824C3110 062CDB50 C1D7D9D7 06366010 | b<.....&APRP...-
+0020 C5E7C9E3 1F000018 06337250 063552D8 | EXIT.....&...Q
+0030 825BC9F6 80000000 E3E2C9D9 06366010 | b$I6....TSIR...-
+0040 C4E2D740 1F000110 06337250 063372A0 | DSP.....&....
+0050 062CDB50 062CDB50 C1D7D9D7 06366010 | ...&...&APRP...-
+0060 D8E4C500 1F4B2110 06337250 063372C0 | QUE.....&...{
+0070 81EAC990 062CDB50 C1D7E4C5 06366010 | a.I....&APUE...-
+0080 C5E7C9E3 1F000050 06337250 063372A0 | EXIT...&...&....
+0090 81EAC9E8 80000000 C1D7D9D7 06366010 | a.IY....APRP...-
+00A0 D9C5D3E2 1F170000 06337250 06366010 | RELS.....&...-
+00B0 824CD2A6 00000000 06366010 00000000 | b<Kw.....-.....
+00C0 E2D9C2E7 1F000000 06337250 80000000 | SRBX.....&....
+00D0 00000000 00F05500 00F05500 824C5BD8 | .....0...0..b<$Q
+00E0 E2D9C2C4 1F000000 06337250 80000000 | SRBD.....&....
+00F0 00000000 00000000 007F0D18 00800000 | ..... " .....
+0100 C9D5E3E7 1D050000 F0C2C6F0 060B0818 | INTX....0BF0....
```

VTVRBLK

VTVRBLK looks at VRs for all subareas and displays the following information:

- Number of subareas supported
- Number of VRBLKs processed
- Number of subareas containing virtual routes
- Number of subareas with no virtual routes
- Number of blocked routes found
- Number of held routes found

For each blocked or held route found, the status areas for each transmission priority (TP0, TP1, and TP2) are analyzed and the following information is displayed if present:

- VR number
- Adjacent subarea
- Destination subarea
- Window sizes
- Pacing limit
- Inbound and outbound sequence numbers
- Selected flags
- VR FSM
- Flow control FSM
- Count of TSCBs on VR hold queue (if any)
- Last pacing request number

Operands

Subarea


Specify a 1–8 hexadecimal digit number in the form X'x...' or a number in the range of 1–2 147 483 647. This represents the number of the subarea. See field ATCSASUP in the ATCVT for the maximum number of subareas available in a dump. Subarea 0 is not valid.

If you are not sure which subarea, if any, is having a problem, specify no subarea to analyze all VRs for all destination subareas.

Report data

Use **Detail** to display further information on every VR found, similar to the information described above for blocked and held routes. **Scan** is the default and provides count information based on the search criteria.

Use the following syntax as an alternative to the panel interface.

Syntax**Additional information**

See z/OS Communications Server: SNA Data Areas Volume 1 for more information on the VRBLK and its contents.

Sample output

VTVRBLK

VTVRBLK Analysis

Route Blocked

```

Subarea (Decimal) =          1

VR block 00690F18      VR number 02      ADJSUBA 00000003      DESTSUBA 00000001
Transmission priority 0  VRFSM          05  FCFSM          00
Window sizes:
  Current (VRBPALIM)           06
  Minimum (VRBMINWS)          02
  Maximum (VRBMAXWS)           06
  Pacing request send count    0000
  Inbound sequence number      0218
  Outbound sequence number     020E
  Last pacing for request number 0212
Status flags:
VRBCWRI = 0  Route change window response NOT required
VRBRWI  = 0  Route RESET window NOT required
VRBHLD  = 1  Half session held NOT required
VRBSCNHQ = 0  All HSQHs have been checked for held sessions

Host subarea skipped          2

VRBLKs processed              32
Subareas supported            511
Subareas with VRs             8
Subareas without VRs          503
Blocked routes                 1
Held routes                     0
  
```

VTWRE

Use VTWRE to count or help analyze waiting request elements (WREs). A WRE represents a VTAM process that is waiting for the completion of some event. A WRE contains a pointer to an event ID that indicates the reason for the wait state. WREs are queued to locked queue anchor blocks (LQABs). The LQABs to be looked at might be:

- All LQABs
- An SSCP I/O LQAB for a specific subarea
- Another specific LQAB

If you invoke VTWRE without operands, the number of WREs queued to all LQABs are counted and the counts are displayed.

Use the DETAIL option to get additional information on each WRE. Also, several operands are available to limit processing to a specific LQAB or to specific WREs.

Operands

Event ID

Enter a 1–100 hexadecimal digit value to be used in matching a WRE event ID found in the dump. The specified event ID is left-aligned when comparing with the contents of the dump. For example, if X'1234' is specified and the dump contained X'F1F0F41234', this would not be a match.

Event ID code

Enter a 4 hexadecimal digit value. If the entered code is fewer than 4 digits, results are unpredictable. Only WREs containing this event ID code are processed. The event ID code identifies the reason for the wait state.

LQAB

Enter a specific LQAB name from the following list to limit processing to a single LQAB. The following LQABs, which are pointed to by the ATCVT, can be examined. The default is ALL.

LQAB name

Description

IOLQB

SSCP I/O LQABs (one per attached subarea)

LUSMQ

Service manager LQAB

MCQAB

Miscellaneous command LQAB

PULQB

Physical unit services LQAB

NODAT_CPWAIT_QUEUE

Network operator services LQAB

SSLQB

Miscellaneous LQAB for session services

SSMQB

Second miscellaneous LQAB for session services

All

All of the above LQABs

If a subarea is specified, the LQAB must be entered as IOLQB. If ALL is specified, all subareas defined to IOLQB by ATCSASUP will be processed.

Subarea

Use this operand to limit processing to the SSCP I/O LQAB for a specific subarea. When a subarea is specified, IOLQB must be specified for the Queue, and Subarea is used as an index into the SSCP I/O LQABs. Specify a 1–8 hexadecimal digit number in the form X'x...'. or a number in the range of 1–2 147 483 647. This represents the number of the subarea. See field ATCSASUP in the ATCVT for the maximum number of subareas available in a dump. Subarea 0 is not valid.

Mask

Enter a 1–100 hexadecimal digit mask. The mask is left-aligned and ANDed with the event identifier in the dump to determine whether the specified event identifier was found.

Note: Mask must be used with Event ID.

Max

Enter a number in the range of 1–99 999 (1–5 decimal digits or 1–4 hexadecimal digits) for the maximum number of WREs to be processed for the selected LQABs. The default is 100.

Control op code

Only WREs containing this CPCB op code are eligible for selection. The control op code must be 1–8 hexadecimal digits in the form X'x...'. If the op code is fewer than eight digits, it is left-aligned and compared with the leftmost digits in the dump.

User correlator

Enter a 1–8 character value. Only WREs containing this user request correlator (URC) are processed. The URC is typically the resource name of the target of a request.

Format

Use **Format** to format the WRE and the EID, if present. With **Noformat**, which is the default, the WRE and EID are not formatted.

Note: Do not specify **Format** if you use **Detail** for Report Data. **Format** is valid only for **Scan**.

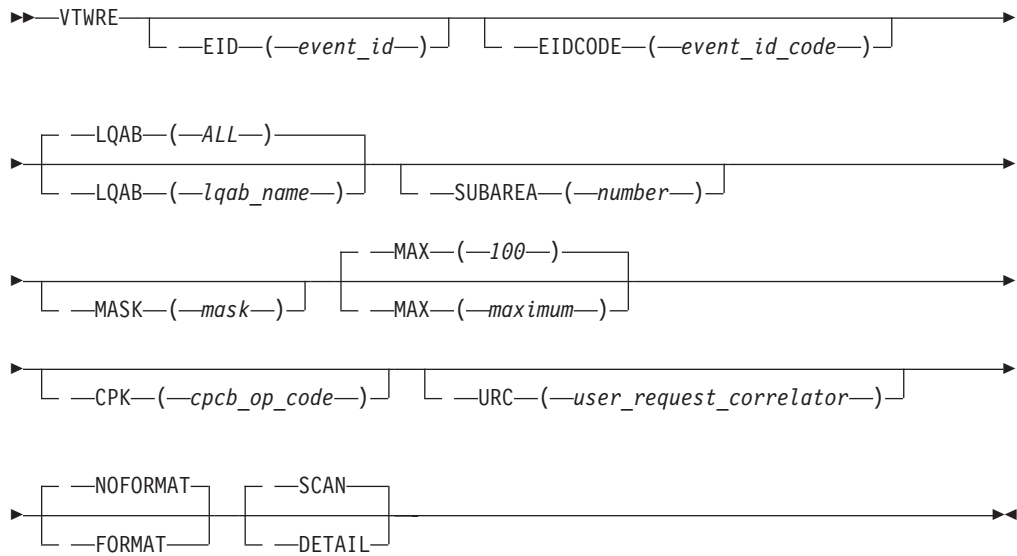
Report data

Scan, which is the default, counts and displays the number of WREs meeting the specified search criteria. Use **Detail** to have the following fields extracted and displayed from each selected WRE:

- WRE address
- Control block (RUPE) address
- CPCBOPC
- URC
- Event ID
- Event ID Code

Use the following syntax as an alternative to the panel interface.

Syntax



Sample output

VTWRE LQAB(IOLQB) DETAIL

```

VTWRE Analysis
ATCIOLQB
WRE ADDR RUPEADDR CPCBOPC   URC   CODE          EVENT ID
-----
0666D200 0664B840 08810680 ECH050Z 0201 000000010071000000010001020108810680
Elements found:           1
Elements processed:      1
  
```

Chapter 7. Using traces

This topic describes when to use traces and shows where in the network you can use each trace to collect data (see Figure 35 on page 308). Examples are included to help you interpret trace output.

This topic includes the following information:

- “Traces provided by VTAM”
- “Traces provided by NCP” on page 353

Traces provided by VTAM

The VTAM program provides several kinds of traces to record the flow of network events. Each trace occurs at a different point in the network (see Figure 35 on page 308). This difference allows you to narrow down the problem by following a request/response unit (RU) through the network and determining where in the network the RU is incorrect. (The RU could be out of sequence or lost, the data in the RU could have been changed, and so on.)

This topic includes the following information:

- “Activating network traces” on page 308
- “Starting the generalized trace facility (GTF)” on page 322
- “Formatting and printing trace records” on page 323
- “Trace output” on page 325
- The APPN route selection trace shows the flow of information throughout the APPN session setup route selection process. See “APPN route selection trace” on page 325 for more information
- VTAM traces and their results:
 - The buffer contents trace shows the contents of inbound and outbound message buffers. See “Buffer contents trace” on page 328 for more information.
 - The I/O trace shows (in order) all I/O sent between VTAM and a particular network resource. See “I/O trace” on page 340 for more information.
 - The QDIOSYNC trace is used to synchronize host and OSA-Express2 or later diagnostic data. See “QDIOSYNC trace” on page 341 for more information.
 - The resource state trace creates VTAM internal trace (VIT) entries when the current state or desired state, or both, of a resource for which tracing has been requested changes. See “Resource state trace” on page 345 for more information.
 - The session management exit (SME) buffer trace shows the input and output of the session management exit (SME) ISTECAA. See “Session management exit (SME) buffer trace” on page 346 for more information.
 - The SMS (buffer use) trace shows information about the use of buffers, including how often a buffer pool has expanded, how many buffers are currently being used, and what was the maximum number of buffers used since the last trace record was written. See “SMS (buffer use) trace” on page 348 for more information.

- The TGET/TPUT trace shows each message as it passes between a TSO command processor and TSO/VTAM. See "TGET/TPUT trace for TSO/VTAM" on page 350 for more information.

The VTAM internal trace (VIT) is discussed in z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT.

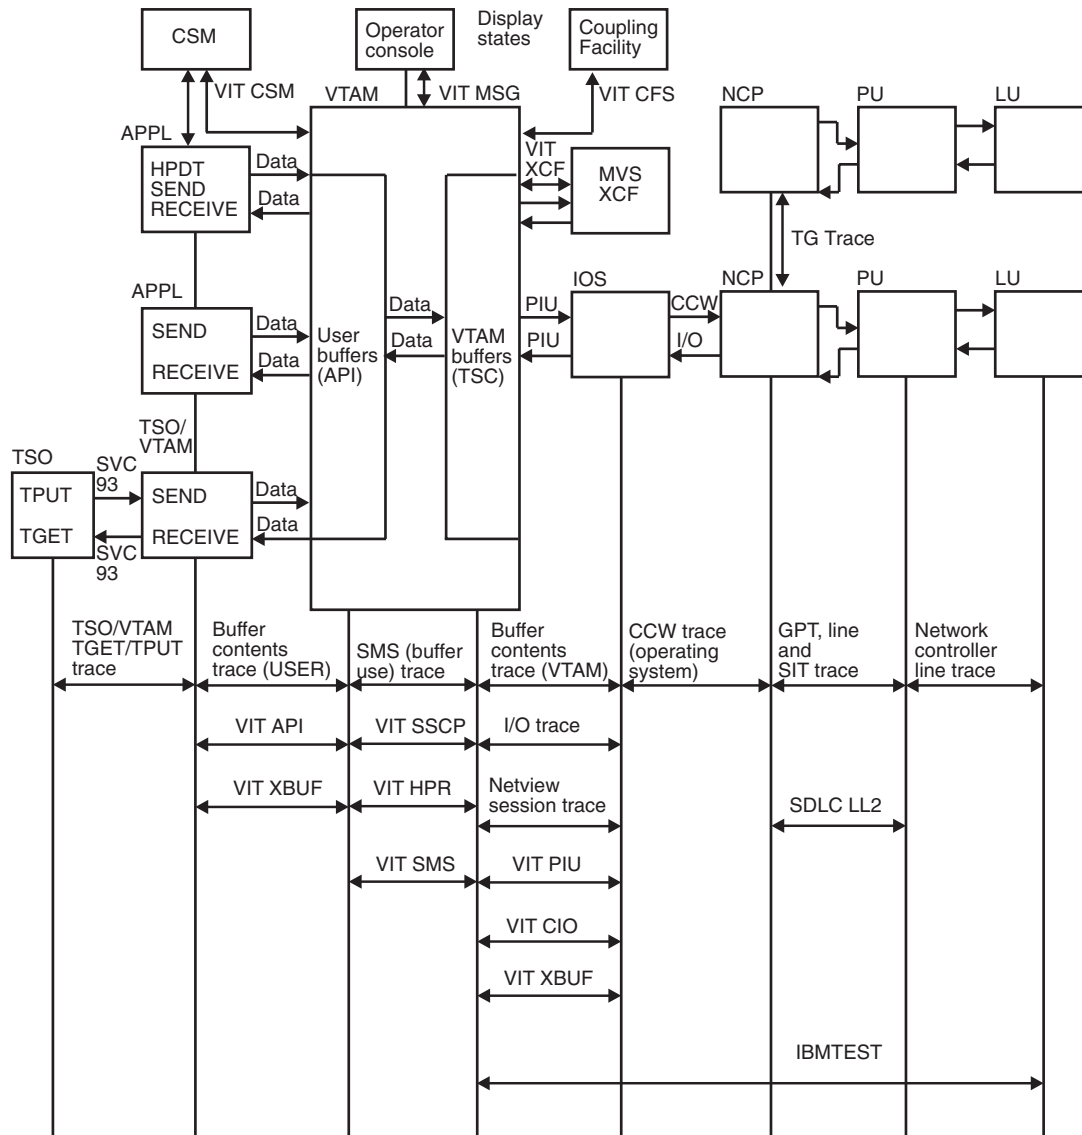


Figure 35. Network traces provided by VTAM

Activating network traces

You can activate VTAM traces when you start VTAM, using the TRACE option on the START command, or you can activate them when VTAM is already running, using the MODIFY TRACE command. This information shows the format of both commands for each type of trace.

Rules:


- The TRACE start option and its qualifiers must be coded on one line.
- The GTF *must* be active to record VTAM traces externally.

Note: VTAM Internal Traces wrap. VTAM External Traces may or may not wrap, depending on the media and GTF specifications.

For more information on activating these traces and optional operands not shown in this table, see *z/OS Communications Server: SNA Resource Definition Reference* and *z/OS Communications Server: SNA Operation*.

APPN route selection trace


Cannot be activated with the TRACE start option.



Buffer contents trace

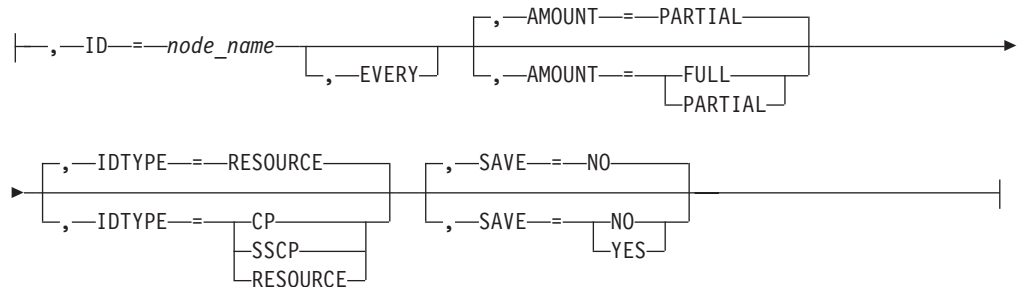


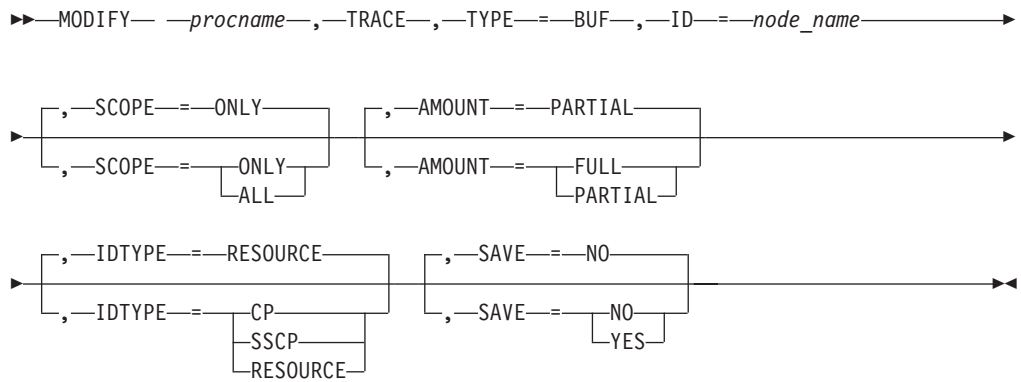
Options:



Notes:

- 1 Do not use NOTRACE when starting VTAM, except to override a TRACE start option coded in a predefined list.
- 2 Code TRACE and its qualifiers on one line. Code the TYPE qualifier immediately following TRACE.





CCW trace



See your operating system books for more information on the CCW trace.

Note: If you have an HPDT MPC connection, you must specify the PCI option when running the CCW trace.


Note: CCW trace will not capture data for a data device for the following devices:

- OSA-Express QDIO
- HiperSockets

I/O trace must be used for these devices. CCW trace can be used for the control devices for the above devices.


CNM trace (communication network management)

Cannot be activated with the TRACE start option.



GPT (generalized PIU trace)

Cannot be activated with the TRACE start option.







I/O trace

►► **START** *procname* , , , (*Options*)


Options:



Notes:

- 1 Do not use NOTRACE when starting VTAM, except to override a TRACE start option coded in a predefined list.
- 2 Code TRACE and its qualifiers on one line. Code the TYPE qualifier immediately following TRACE.
- 3 SAVE=YES is the default if issued on the START command, and SAVE=NO is the default if issued via MODIFY TRACE.

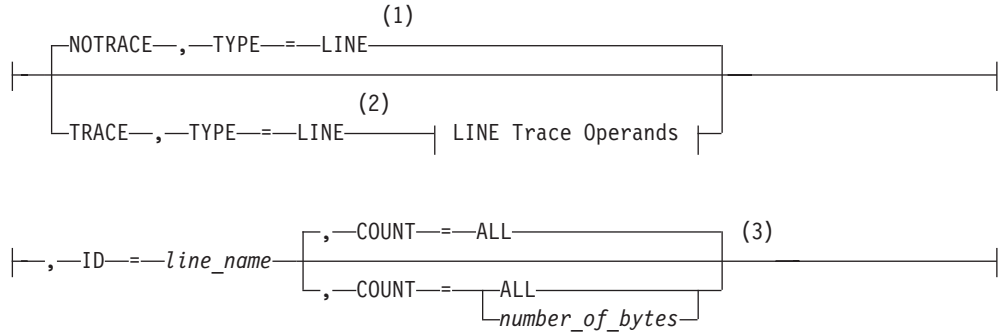
►► **MODIFY** *procname* , TRACE, TYPE=IO, ID=*node_name*



Line trace

►► START — *procname* —, —, —, — (— | Options | —) —————►►


Options:



Notes:

- 1 Do not use NOTRACE when starting VTAM, except to override a TRACE start option coded in a predefined list.
- 2 Code TRACE and its qualifiers on one line. Code the TYPE qualifier immediately following TRACE.
- 3 COUNT applies only to the IBM 3720 and 3745 communication controllers.


►► MODIFY — *procname* —, —TRACE—, —TYPE—, —LINE—, —ID—, —line_name— —————►►

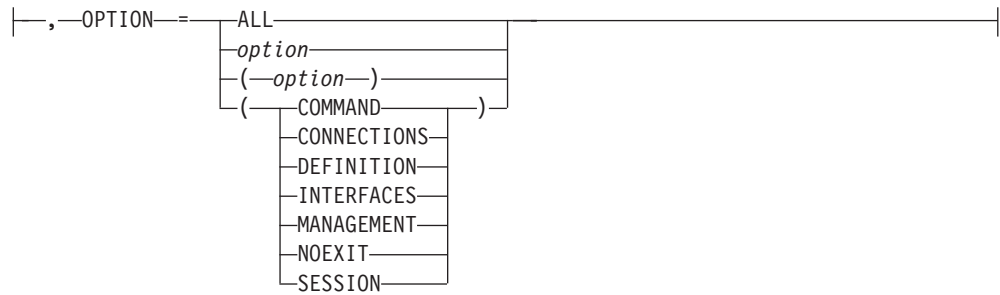


Module Trace

►► START — *procname* —, —, —, — (— | Options | —) —————►►

Options:

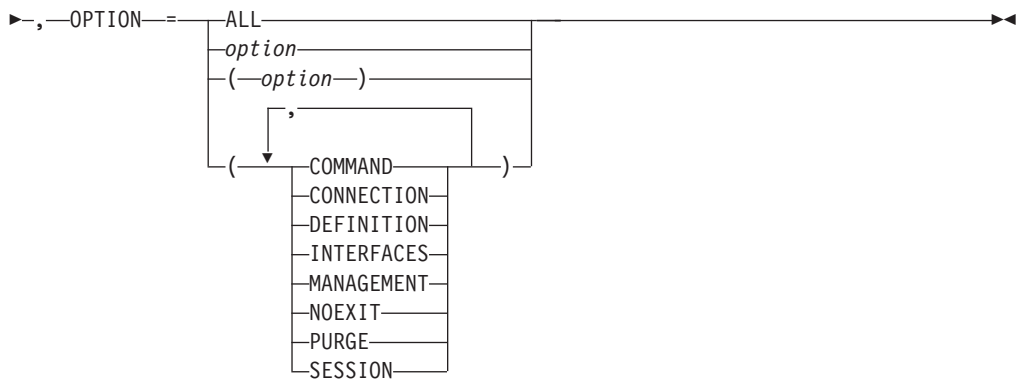




Notes:

- 1 Do not use NOTRACE when starting VTAM, except to override a TRACE start option coded in a predefined list.
- 2 Code TRACE and its qualifiers on one line. Code the TYPE qualifier immediately following TRACE.

►► MODIFY — *procname* —, —TRACE—, —TYPE—= —MODULE—►►



NetView session trace


►► STARTCNM — NLDM — TRACE ►►

For more information on the NetView session trace, see *Tivoli® NetView for z/OS Version 5.2 Command Reference Volumes 1 & 2*.

Network controller line trace

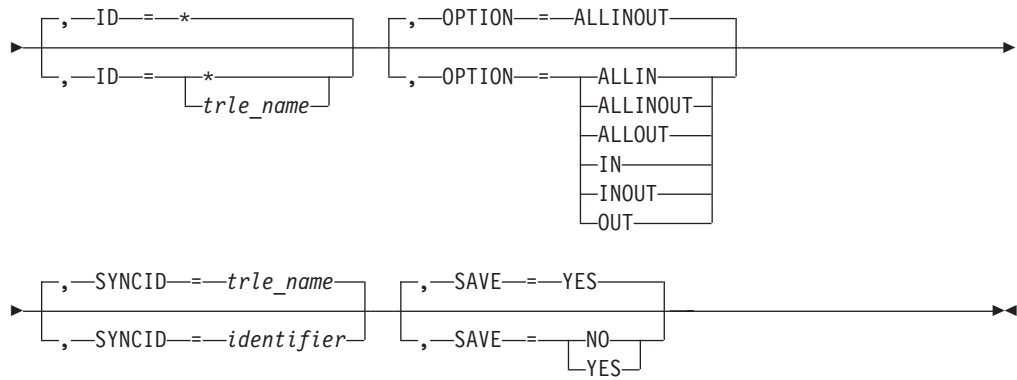
Cannot be activated with the TRACE start option.

►► MODIFY — *procname* —, —TRACE—, —TYPE—= —NETCTLR—, —ID—= —*pu_name*—►►



QDIOSYNC trace


►► MODIFY — *procname* —, —TRACE—, —TYPE—= —QDIOSYNC—►►



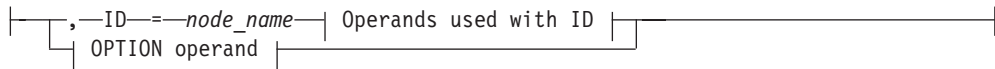
Resource state trace

►► START —procname—, —, —, — (— Options —) —►►

Options:



STATE trace operands:

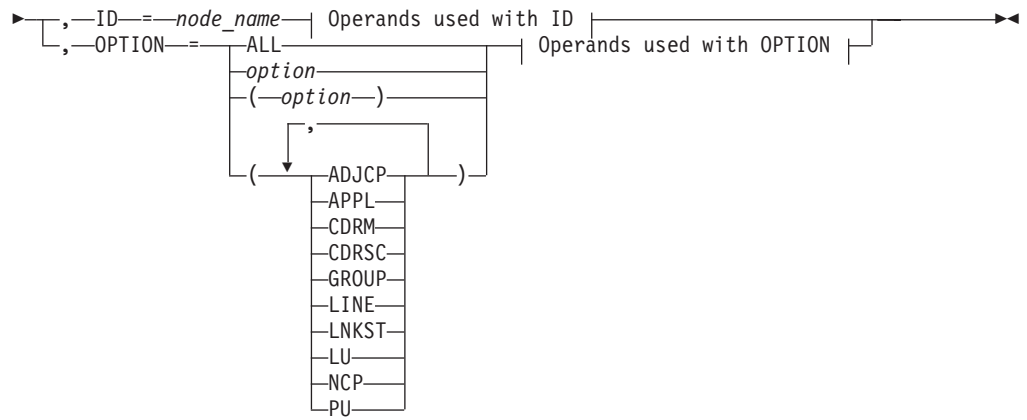


Notes:

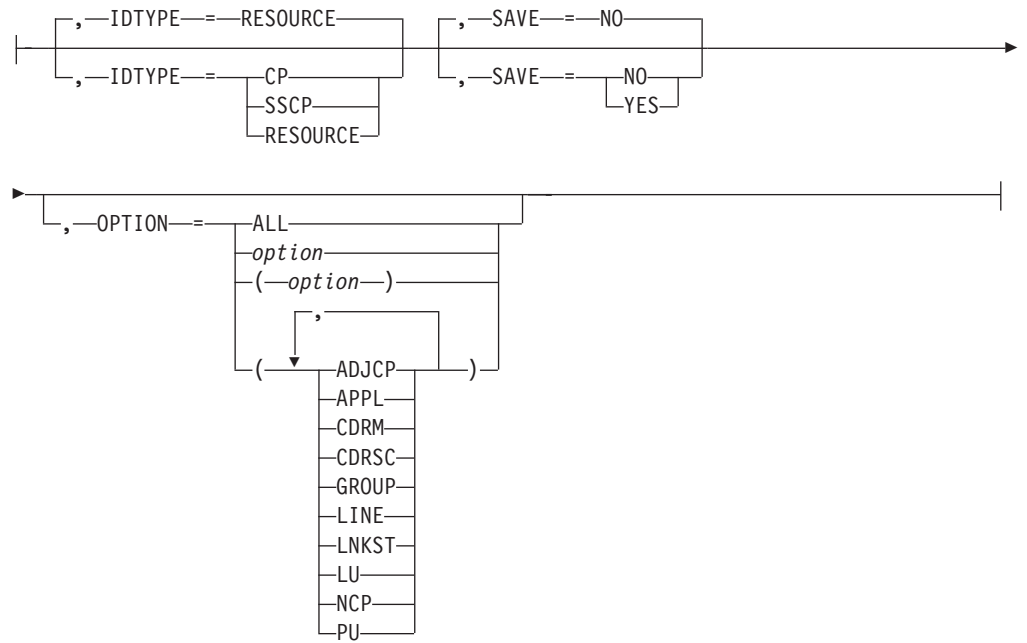
- 1 Do not use NOTRACE when starting VTAM, except to override a TRACE start option coded in a predefined list.
- 2 Code TRACE and its qualifiers on one line. Code the TYPE qualifier immediately following TRACE.

Resource state trace (continued)

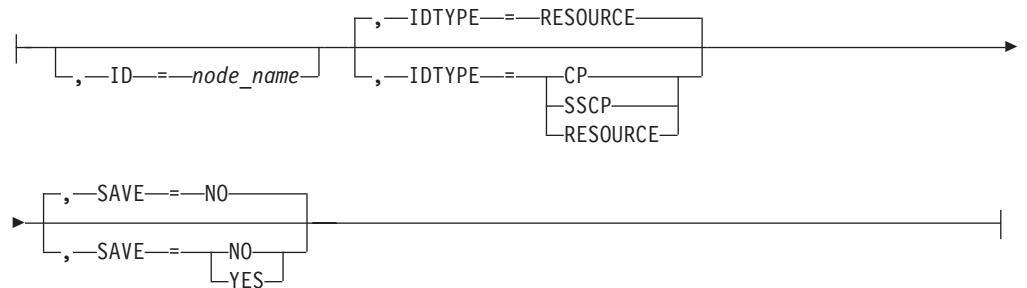
►► MODIFY —procname—, —TRACE—, —TYPE=STATE —►►



Operands used with ID:



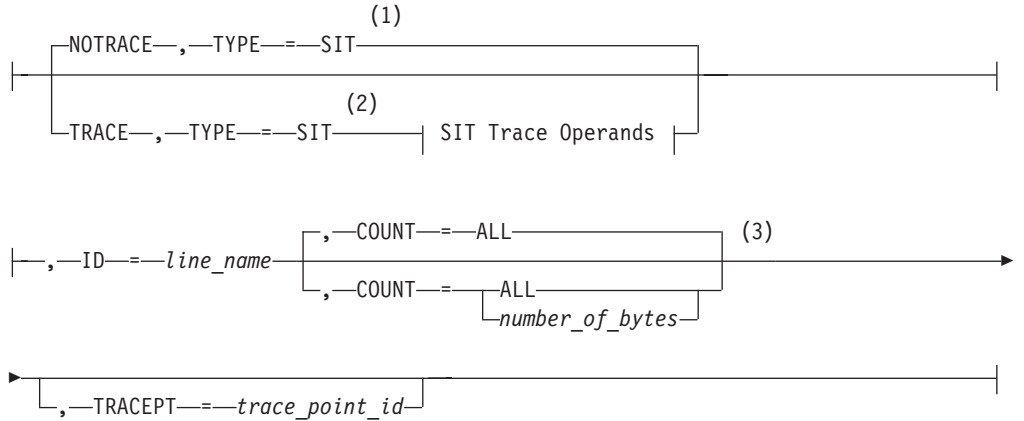
Operands used with OPTION:



SIT (scanner interface trace)

►► START *—procname—*, —, —, — (— | Options | —) ►►


Options:



Notes:

- 1 Do not use NOTRACE when starting VTAM, except to override a TRACE start option coded in a predefined list.
- 2 Code TRACE and its qualifiers on one line. Code the TYPE qualifier immediately following TRACE.
- 3 COUNT applies only to the IBM 3720 and 3745 communication controllers.

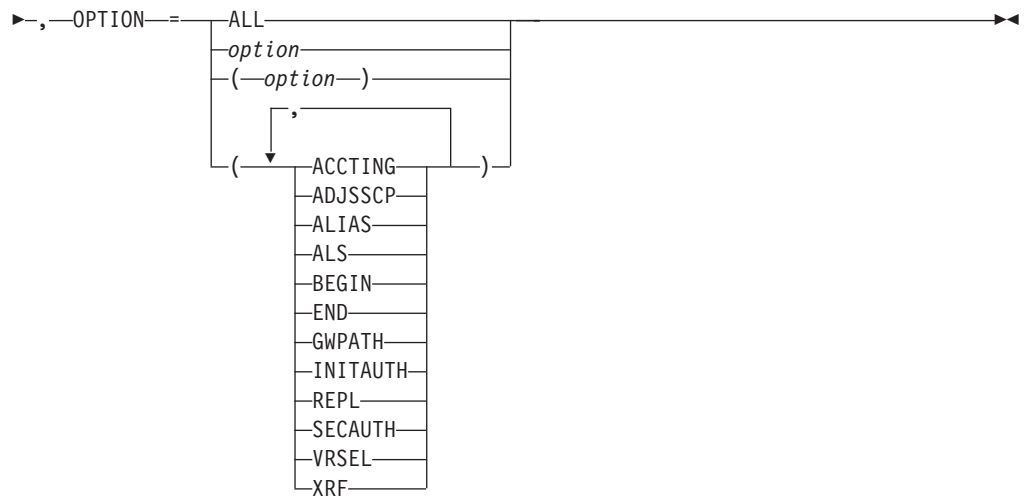
►► MODIFY *—procname—*, —TRACE—, —TYPE—, —SIT—, —ID—, —line_name— ►►



SME Buffer trace (Session Management Exit)

Cannot be activated with the TRACE start option.


►► MODIFY *—procname—*, —TRACE—, —TYPE—, —EXIT—, —ID—, —ISTEXCAA— ►►



SMS (buffer use) trace

▶▶ START *procname*, , , , (Options) ▶▶

Options:



Notes:

- 1 Do not use NOTRACE when starting VTAM, except to override a TRACE start option coded in a predefined list.
- 2 Code TRACE and its qualifiers on one line. Code the TYPE qualifier immediately following TRACE.

TG trace (transmission group)

Cannot be activated with the TRACE start option.

▶▶ MODIFY *procname*, TRACE, TYPE=TG, ID=*line_name* ▶▶

TSO/VTAM TGET/TPUT trace

Cannot be activated with the TRACE start option.

▶▶ MODIFY *procname*, TRACE, TYPE=TSO, ID=*tso_user_id* ▶▶

VTAM internal trace (VIT)

▶▶ START *procname*, , , , (Options) ▶▶

Note:

1. Precede the option list with three commas and enclose the group of options in parentheses.
2. Start options that are entered on the START command must be separated by commas. Do not leave any blanks between options.


Options:

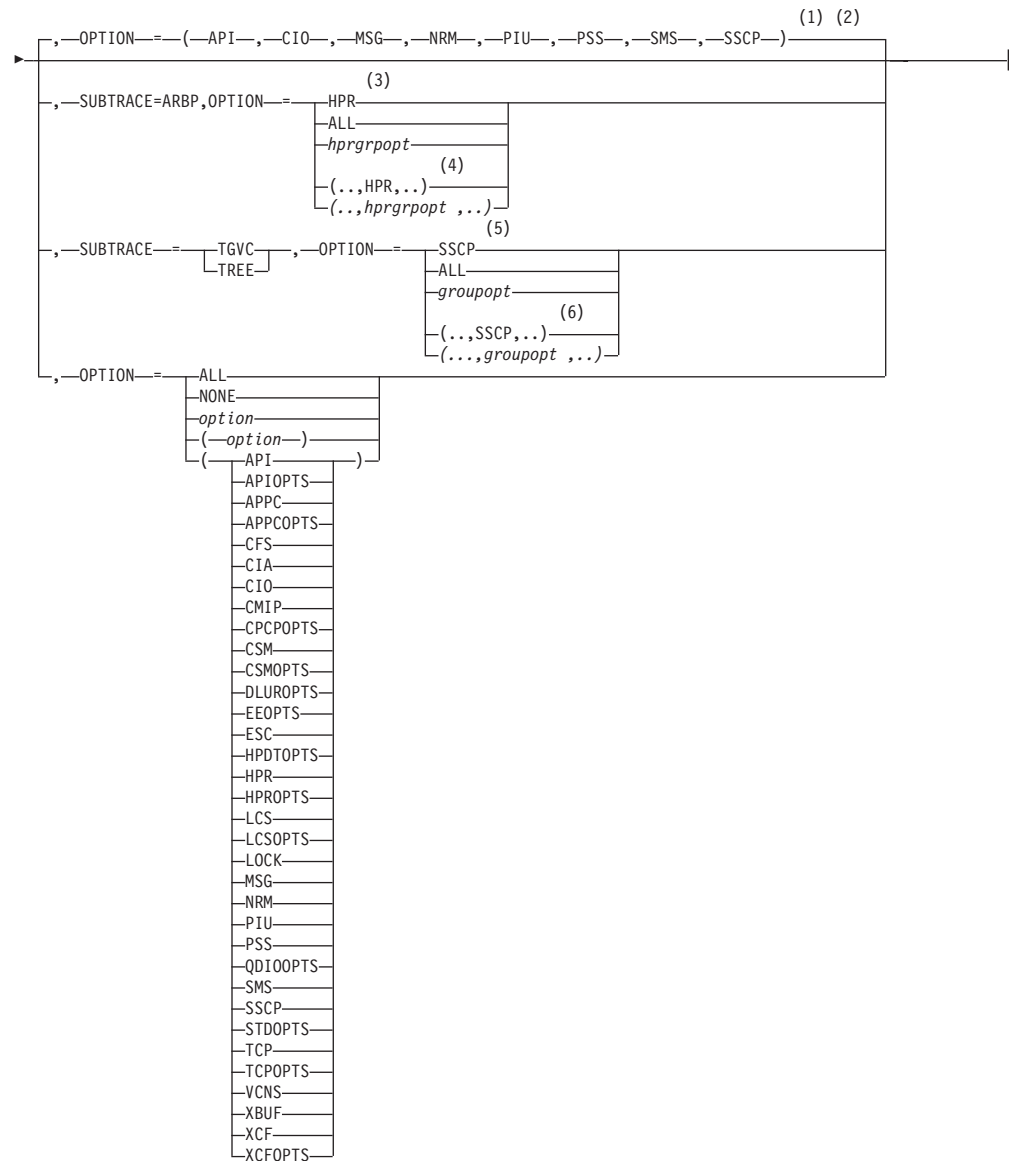


Notes:

- 1 Code TRACE and its qualifiers on one line.
- 2 NOTRACE,TYPE=VTAM is accepted but ignored. Tracing is started with the default trace table size and the default options.

VIT operands:




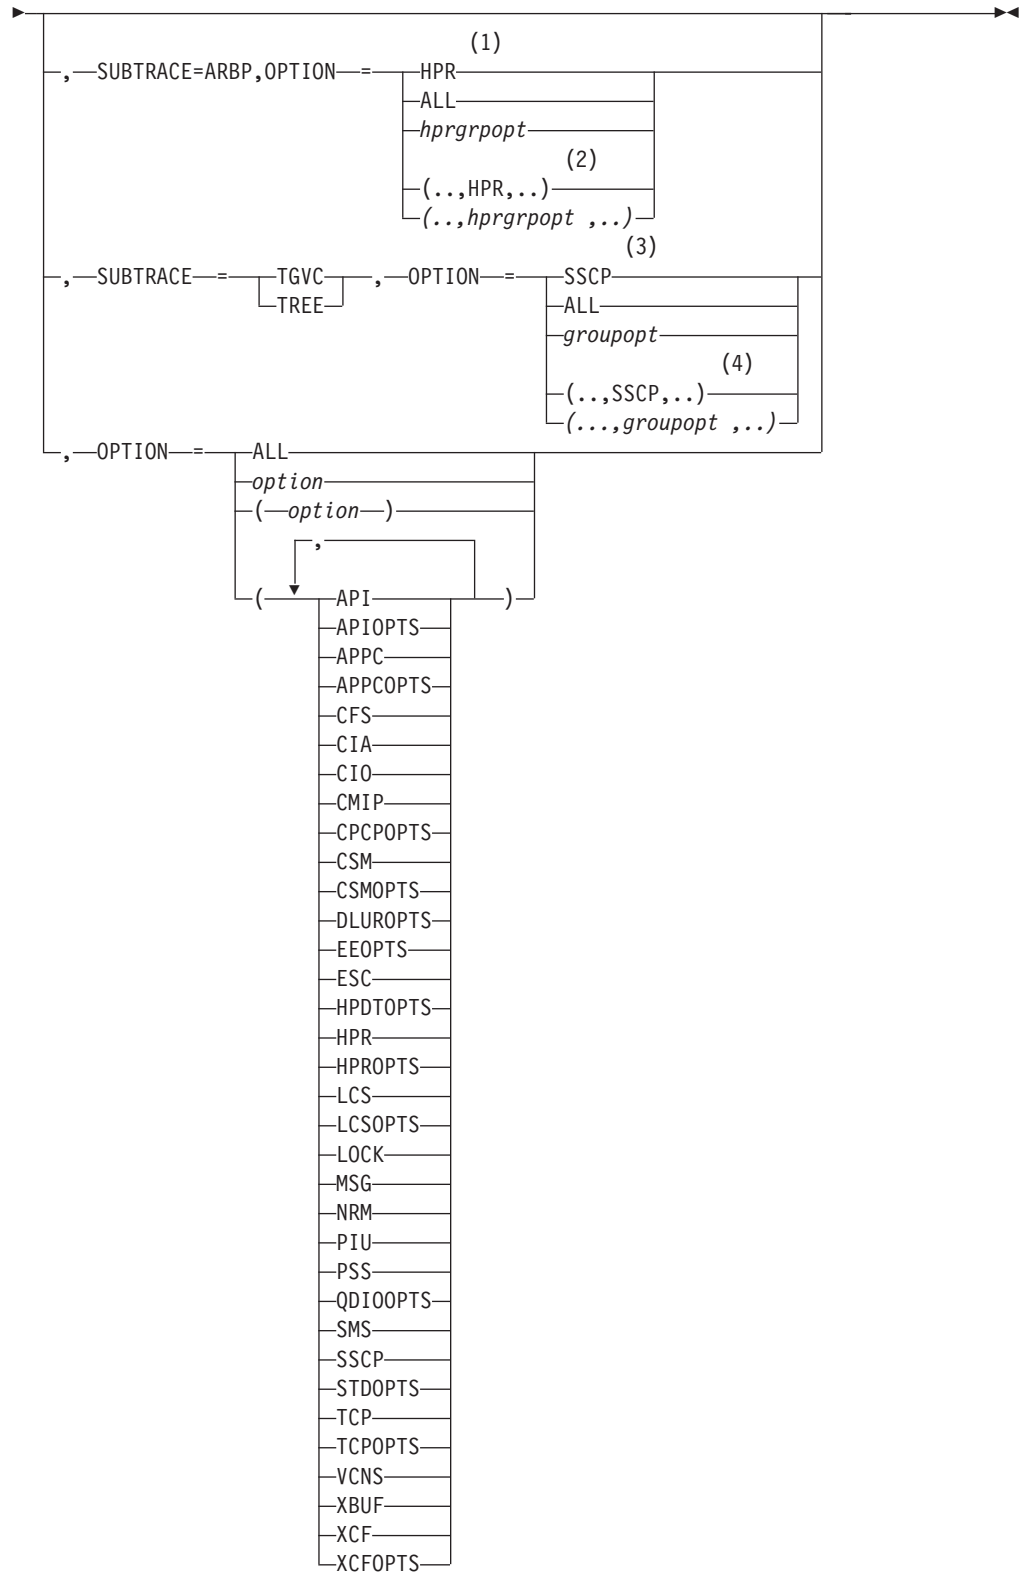


Notes:

- 1 The default options apply only to `MODE=INT`.
- 2 You can turn off `PSS` and `SMS`.
- 3 When `SUBTRACE=ARBP` is specified, if a single `OPTION` value is coded, it must be `HPR`, `ALL`, or one of the group options (*hprgrpopt*) that include `HPR` as an individual option equivalent. The applicable group options are `DLUROPTS`, `EEOPTS`, `HPDTPPTS`, `HPROPTS`, `QDIOOPTS`, and `XCFOPTS`.
- 4 If multiple trace options are coded in parentheses, either `HPR` or one of the group options (*hprgrpopt*) that include `HPR` as an individual option equivalent must be coded inside the parentheses when `SUBTRACE=ARBP` is coded.
- 5 When `SUBTRACE=TGVC` or `SUBTRACE=TREE` is coded, if a single `OPTION` value is coded, it must be `SSCP`, `ALL`, or one of the group options (*groupopt*), all of which include `SSCP` as an individual option equivalent. The group options are `APIOPTS`, `APPCOPTS`, `CPCPOPTS`, `CSMOPTS`, `DLUROPTS`, `EEOPTS`, `HPDTPPTS`, `HPROPTS`, `LCSOPTS`, `QDIOOPTS`, `STDOPTS`, `TCPOPTS`, and `XCFOPTS`.

- 6 If multiple trace options are coded in parentheses, either SSCP or one of the group options (*groupopt*) must be coded inside the parentheses when SUBTRACE=TGVC or SUBTRACE=TREE is coded.





Notes:

- 1 When `SUBTRACE=ARBP` is specified, if a single `OPTION` value is coded, it must be `HPR`, `ALL`, or one of the group options (`hprgrpopt`) that include `HPR`

as an individual option equivalent. The applicable group options are DLUROPTS, EEOPTS, HPDTPPTS, HPROPTS, QDIOOPTS, and XCFOPTS.

- 2 If multiple trace options are coded in parentheses and SUBTRACE=ARBP is coded, then either HPR or one of the group options (*hprgrpopt*) that include HPR as an individual option equivalent must be coded inside the parentheses.
- 3 When SUBTRACE=TGVC or SUBTRACE=TREE is coded, if a single OPTION value is coded, it must be either SSCP, ALL, or one of the group options (*groupopt*), all of which include SSCP as an individual option equivalent. The group options are APIOPTS, APPCOPTS, CPCOPTS, CSMOPTS, DLUROPTS, EEOPTS, HPDTPPTS, HPROPTS, LCSOPTS, QDIOOPTS, STDOPTS, TCPOPTS, and XCFOPTS.
- 4 If multiple trace options are coded in parentheses, either SSCP or one of the group options (*groupopt*) must be coded inside the parentheses when SUBTRACE=TGVC or SUBTRACE=TREE is coded.

Starting the generalized trace facility (GTF)

Because VTAM passes all external trace data to the generalized trace facility (GTF), the GTF must be active to use VTAM traces. Specify TRACE=USRP to receive a prompt for all VTAM traces. You will then be able to select the specific event identifier (EID) to be traced. See APAR II03922 for additional information.

Note: If you do not limit the GTF trace output using the USRP option, the GTF collects all USR events issued in the MVS system, often resulting in a large amount of unwanted information.

I/O trace entries: The external VIT is now used to record the I/O trace entries. PIU, NLPI, NLPO, LSNA, and MPTNFMT entries may be written for a specific I/O trace invocation. For more information, see MODIFY TRACE or MODIFY NOTRACE in z/OS Communications Server: SNA Operation.

```

NC000000 V535M433 93306 14:25:04.88 01 00000290 S FVGT.F.B
N 4020000 V535M433 93306 14:25:10.06 STC00017 00000090 AHL121I TRACE OPTION INPUT INDICATED FROM MEMBER GTFPARM OF PDS
S
N 0020000 V535M433 93306 14:25:10.36 STC00017 00000090 TRACE=SYSM,USR,TRC,DSP,PCI,SRM
N 4020000 V535M433 93306 14:25:10.38 STC00017 00000090 AHL103I TRACE OPTIONS SELECTED --SYSM,USR,TRC,DSP,PCI,SRM
W 4020000 V535M433 93306 14:25:10.38 STC00017 00000090 *25 AHL125A RESPECIFY TRACE OPTIONS OR REPLY U
NC000000 V535M433 93306 14:25:19.47 01 00000290 R 25,TRACE=USRP
NR4020000 V535M433 93306 14:25:19.57 00000090 IEE600I REPLY TO 25 IS;TRACE=USRP
N 0020000 V535M433 93306 14:25:19.73 STC00017 00000090 TRACE=USRP
W 4020000 V535M433 93306 14:25:19.87 STC00017 00000090 *26 AHL101A SPECIFY TRACE EVENT KEYWORDS --USR=
NC000000 V535M433 93306 14:25:35.68 01 00000290 R 26,USR=(FEF,FF1)
NR4020000 V535M433 93306 14:25:35.76 00000090 IEE600I REPLY TO 26 IS;USR=(FEF,FF1)
N 0020000 V535M433 93306 14:25:35.88 STC00017 00000090 USR=(FEF,FF1)
W 4020000 V535M433 93306 14:25:35.94 STC00017 00000090 *27 AHL102A CONTINUE TRACE DEFINITION OR REPLY END
NC000000 V535M433 93306 14:25:41.00 01 00000290 R 27,END
NR4020000 V535M433 93306 14:25:41.09 00000090 IEE600I REPLY TO 27 IS;END
N 0020000 V535M433 93306 14:25:41.25 STC00017 00000090 END
N 4020000 V535M433 93306 14:25:41.25 STC00017 00000090 AHL103I TRACE OPTIONS SELECTED --USR=(FEF,FF1)
W 4020000 V535M433 93306 14:25:41.25 STC00017 00000090 *28 AHL125A RESPECIFY TRACE OPTIONS OR REPLY U
NC000000 V535M433 93306 14:25:46.83 01 00000290 R 28,U
NR4020000 V535M433 93306 14:25:46.99 00000090 IEE600I REPLY TO 28 IS;U
N 4020000 V535M433 93306 14:25:48.30 STC00017 00000090 AHL031I GTF INITIALIZATION COMPLETE

```

Figure 36. Starting GTF for the VTAM buffer contents trace

See Table 48 on page 649 to determine which document contains more information on the GTF.

Formatting and printing trace records

Table 13 indicates which traces can be formatted and printed by each of the formatting programs.² Descriptions of the programs appear later in the topic. In this table:

- FP indicates that the trace is both formatted and printed (for VTAM records only).
- UP indicates that the trace is printed but not formatted.
- A blank entry indicates that the trace is not formatted or printed.

Table 13. Processing externally recorded trace data

Trace	ACF/TAP	IPCS
Buffer contents trace	FP	FP
Generalized PIU trace	FP	
I/O trace	UP	FP
Line trace ²	FP	FP ¹
Network controller line trace	FP	FP ¹
Scanner interface trace	FP	
SME buffer trace		FP
SMS (buffer use) trace	UP	FP
TGET/TPUT trace	UP	UP
Transmission group trace	FP	FP ¹
VTAM internal trace ³	UP	FP

Note:

1. Only scanner type 1, 2, and 3 records are processed. All others must be processed using ACF/TAP.
2. ACF/TAP must be used except for 3705 traces. See “Line trace operation” on page 355 for more information.
3. ACF/TAP allows you to specify some formatting parameters if the VIT is running in external mode.

Using ACF/TAP

Use the trace analysis program (ACF/TAP) to print all VTAM external trace entries or up to ten entry types. Table 14 lists the options to use on the INPUT operand for formatting and printing traces. See *z/OS Communications Server: ACF/TAP Trace Analysis Handbook* for more information.

Table 14. Printing external trace entries

Specify:	To format and print:
LINE	Line trace
LINE	Scanner interface trace (3720, 3725, and 3745 only)
LINE	Transmission group trace

². This is not a complete list of programs that process external trace data.

Table 14. Printing external trace entries (continued)

Specify:	To format and print:
GPT	Generalized PIU trace
BUFFER	Buffer contents trace
NETCTLR	Network controller line trace

Notes on using ACF/TAP:

1. You can use ACF/TAP to format the VTAM full buffer contents trace. Only the first 256 bytes are traced.
2. ACF/TAP does not print a buffer contents trace that is traced at the API. Use IPCS GTFTRACE to print API (FF1) traces.

Using IPCS with the GTF trace option

To format and print VTAM traces, set *USR(symnum 1[,symnum2]...[,symnum6])|ALL* on the GTFTRACE option.

For *symnum*, use either a symbolic name or a number representing the trace that you want formatted and printed. If you specify *USR(ALL)*, IPCS formats and prints all user and subsystem traces recorded by the GRF. For information on starting the GRF, see “Starting the generalized trace facility (GTF)” on page 322.

Table 15 lists the valid symbols and numbers for the VTAM traces.

Table 15. Symbols and numbers for formatting and printing VTAM traces

Symbol	Number	Trace
INT1 ³	FE1	VTAM internal trace, I/O trace
TPIO	FEF	VTAM buffer contents trace (TSC component) Trace output says “VTAM.”
CL01	FF1	SME buffer trace
CL01	FF1	VTAM buffer contents trace (API component) Trace output says “USER.”
CL02	FF0	SMS (buffer use) trace
LINE	FF2	NCP 37xx line or TG trace
APTH	FE2	TSO/VTAM TGET/TPUT trace
APTD	FE4	Line PIU, generalized PIU, or network controller line trace

Note: The symbol and the number can be used interchangeably, for example, *USR(LINE)* or *USR(FF2)*; however, when starting the GTF, use the number.

See Table 48 on page 649 to determine what document describes how to use the GTF and IPCS.

3. I/O trace is now done using VIT.

Trace output

In addition to the fields produced by VTAM, Table 16 contains generic fields that might appear in VTAM trace output.

Table 16. Fields in VTAM trace output

Field header	Meaning
ASCB <i>nnnnnnnn</i>	The address of the ASCB for the address space that created the record.
CPU <i>nnnn</i>	The ID of the host processor in which the trace was run (applies only in a multiprocessor configuration).
JOBN <i>ccccccc</i>	The name of the job associated with the I/O operation (for an I/O trace).

VTAM trace record formats

This information shows the trace record formats for the following traces:

- APPN route selection trace
- Buffer contents trace for the VTAM application programming interface (API) and the transmission subsystem component (TSC)
- Buffer contents trace for the Common Management Information Protocol (CMIP) services management information base (MIB) API
- Line trace

APPN route selection trace

The APPN route selection trace can impact VTAM performance during session setup route selection; it should be activated only when attempting to document a route selection problem, usually when requested by the z/OS Communications Server service organization.

APPN route selection trace operation:

Start the APPN route selection trace with the `MODIFY TRACE,TYPE=ROUTE` command and stop the trace with the `MODIFY NOTRACE,TYPE=ROUTE` command. The status of the APPN route selection trace and the amount of storage allocated to the trace table can be displayed with the `DISPLAY TRACES,TYPE=ROUTE` command. For more information about the `MODIFY TRACE`, the `MODIFY NOTRACE`, and the `DISPLAY TRACES` commands, see *z/OS Communications Server: SNA Operation*.

Gathering documentation:

Because the purpose of the APPN route selection trace is to capture documentation to solve APPN routing problems, use the following procedure to gather the documentation.

Procedure

Perform the following procedure to gather the documentation:

1. Start the APPN route selection trace with the `MODIFY TRACE,TYPE=ROUTE` command.
2. Re-create the problem of the incorrect APPN route being taken.

3. Stop the APPN route selection trace with the MODIFY NOTRACE,TYPE=ROUTE command.
4. Dump VTAM to capture the information in the APPN route selection trace.
5. Free the route selection trace table storage with the MODIFY NOTRACE,TYPE=ROUTE,FREE=YES.

APPN route selection trace output:

The APPN route selection trace is an internal trace table and the trace output is not documented.

Buffer contents trace for VTAM API and TSC

Figure 37 shows the trace record format for the buffer contents trace for the VTAM application programming interface (API) and the transmission subsystem component (TSC).

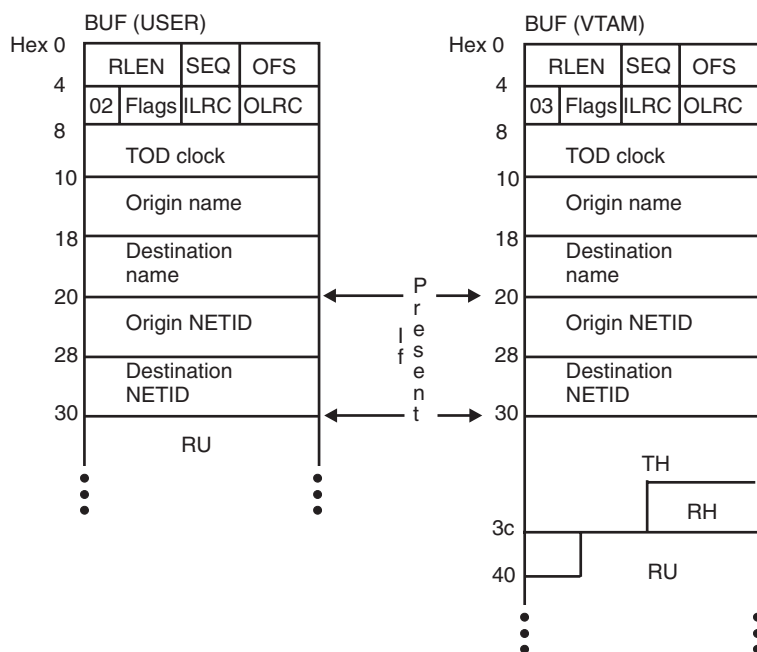


Figure 37. Format for buffer contents trace records for VTAM API and TSC

Buffer contents trace for CMIP services API

Figure 38 on page 327 shows the trace record format for the buffer contents trace for the CMIP services management information base (MIB) API.

The buffer trace records will consist of the standard user buffer trace header (ISTTRAB), but the data portion will contain the CMIP data instead of an RU. The record format is shown in Figure 38 on page 327.

Hex Offset L	BUF (USER)			
	RLEN	SEQ	OFS	
	TYP	FLGS	LRC	OLRC
	TOD clock			
	Origin name			
	Destination name			
	Origin NetID			
	Destination NetID			
	ISTAPIHD			
	ISTTLV ISTTLV ••• (if present)			
	String header (if present)			
	CMIPstring			

Legend

RLEN = Total record length
 SEQ = Sequence number when running full buffer trace
 OFS = Offset of the network qualifiers (if present) or zero
 TYP = Record type:
 2 = USER partial
 11 = USER full
 Flag bit 0: 1 = Inbound
 Flag bit 1: 1 = Confidential text (not applicable for CMIP)
 Flag bit 2: 1 = First segmentt
 Flag bit 3: 1 = Last segment
 ILRC = Inbound lost record count
 OLRC = Outbound lost record count
 TOD clock = Time-of-day clock timestamp
 Origin name¹ = Application program or CMIP services name
 Destination name¹ = Application program or CMIP services name
 Origin NetID² = VTAM's NetID
 Destination NetID² = VTAM's NetID
 ISTAPIHD = The API header
 ISTTLV³ = API tag-length-values (if present)
 String header³ = The CMIP string header (if present)
 CMIP string³ = The CMIP string itself (if present)

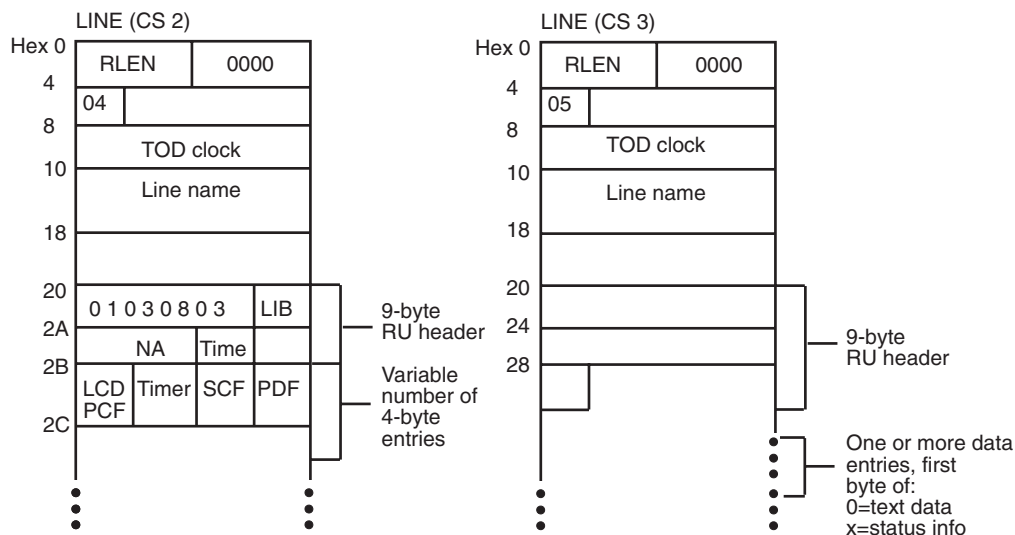
Notes:

1. The names used for origin and destination are the application program name and **ISTCMSVC** for CMIP services.
2. CMIP application programs and CMIP services must have the same (local) NetID.
3. If they are not present, no space is reserved for the ISTTLV, string header, and CMIP string. These are optional portions of the data, dependent on the message that is being sent.

Figure 38. Format for buffer contents trace records for CMIP services API

Line trace

Figure 39 shows the trace record format for line traces.



Legend

RLEN = Total record length

TOD clock = Time-of-day clock timestamp

Figure 39. Format for line trace records

Buffer contents trace

You can run a buffer contents trace for the VTAM API and TSC and for the CMIP services API.

Buffer contents trace for the VTAM API and TSC

The buffer contents trace shows the contents of message buffers in the application programming interface (API) and the transmission subsystem component (TSC). When data is sent by an application program (outbound), API is the first component of VTAM to process it, and TSC is the last component of VTAM to process it. When data is received from the network (inbound), TSC is the first component of VTAM to process it, and API is the last. To enable you to distinguish where in VTAM the trace data was recorded, the output specifies:

- Either USER (for data recorded in API) or VTAM (for data recorded in TSC)
- If the recorded data is inbound or outbound

The API writes user buffer contents trace records while user data is still in the application program's buffers, before it is copied into VTAM buffers. Only user data is recorded.

TSC writes VTAM buffer contents trace records while the data is in VTAM fixed I/O buffers. The PIU (the transmission header - TH, request/response header - RH, and user data) is recorded.

For a list of the resources for which you may request a buffer contents trace, see *z/OS Communications Server: SNA Operation*.

Note:

1. VTAM can start a buffer contents trace only for the resources that it owns. A data host, which does not own any NCPs, cannot start a buffer contents trace for an NCP or any of the NCP's subordinate resources.
2. If you want to trace a session between an LU and an application program, you must start the trace at the host where the application program resides.
3. If the buffer contents trace is active for a specific APPC application, the FMH5 is always traced at the user (API) level.
4. The VTAM TSC component is bypassed for conversation level data exchanged between two VTAM/APPC applications residing on the same host and using the APPCCMD macroinstruction interface to communicate. Thus VTAM buffer contents records are not to be recorded for this data. USER buffer contents records will continue to be recorded.

When to use the buffer contents trace for the VTAM API and TSC:

The buffer contents trace can help you determine whether a problem is in the host (VTAM or an application program) or in the network. For example, if an application program sends a message to a terminal, and the message is correct in VTAM buffer contents trace output, but the message does not appear correctly at the terminal, then the problem is probably in a system resource other than VTAM or the application program.

The buffer contents trace cannot always be used to distinguish an application program problem from a VTAM problem. However, it can confirm the order in which data is passed between an application program and a logical unit. It can also record all the data passing to and from an application program.

If you do not need to trace user data, use the I/O trace.

Buffer contents trace for CMIP services API

The buffer contents trace includes CMIP strings in the user buffer trace. The CMIP services API function that sends the CMIP string is recorded in the VIT entries MDEL, MQRQ, MQRS, MREG, and RQE, all of which are traced under the CMIP VIT option. The invoke identifier field appears in the VIT entries and in the API header portion of the buffer trace data. See *z/OS Communications Server: SNA Data Areas Volume 1* for the APIHDR control block.

When ISTNOTIF is traced, the notifications from ISTTOPAG to ISTNOTIF are traced. However, they are not traced when only ISTTOPAG is traced. Notifications are shown only as they flow from CMIP services to ISTNOTIF. To see notifications, trace ISTNOTIF. To omit notifications, do not trace ISTNOTIF. Note that when you are tracing ISTTOPAG and not ISTNOTIF, notifications might exist that do not appear in the buffer trace.

When to use the buffer contents trace for CMIP services API:

The buffer contents trace can be used to help diagnose problems suspected in CMIP services or the VTAM topology agent. In addition, it can be used to trace data on the CMIP services API for external CMIP application programs, such as the NetView program.

Figure 40 on page 331 illustrates the buffer trace points within CMIP services and the VTAM topology agent.

There are several types of problems that might occur:

1. The manager application program does not receive something it expected to receive.
2. The manager application program receives something it does not expect.
3. The manager application program receives what it expected, but the data does not appear to be correct.
4. CMIP services abends or dumps.
5. The VTAM topology agent abends or dumps.

For each problem, you might see the following symptoms:

For problem 1, you might have activated some resources but never saw anything show up on graphical monitor facility (GMF). If you were monitoring the network by commands, you should start the buffer contents trace with the `ID=ISTTOPAG` and `ID=acbname` operands. If you were monitoring the network implicitly by notifications, for example, by receiving topology updates for definition group objects and the LU object (for the NetView program), or you do not know where the update should be coming from, start the buffer contents trace with the `ID=ISTNOTIF` operand, in addition to the `ID=ISTTOPAG` and `ID=acbname` operands.

For problem 2, the NetView program issues an FFST probe. For information about what values to specify for the ID operand for the buffer contents trace, see the NetView publications. If you cannot determine what buffer trace IDs to run, run the buffer contents trace with the following values for the ID operand:

- `ID=ISTNOTIF`
- `ID=ISTTOPAG`
- `ID=acbname`

For problem 3, in most cases, use the `ISTTOPAG` value for the buffer contents trace. In a few cases, you need to specify the `ISTNOTIF` value for the ID operand.

If you were monitoring the network by commands, you should start the buffer contents trace with the `ID=ISTTOPAG` operand. If you were monitoring the network implicitly by notifications, for example, by receiving topology updates for definition group objects and the LU object (for the NetView program), or you do not know where the update should be coming from, start the buffer contents trace with the `ID=ISTNOTIF` operand.

If you are not sure which value to specify, specify both `ISTTOPAG` and `ISTNOTIF`.

For problem 4, if you can re-create the problem, start the buffer trace with the `ID=acbname` operand that specifies the *acbname* of each CMIP application program using CMIP services on this host. If you are using the VTAM topology agent, specify the `ID=ISTTOPAG` operand as well. Re-create the problem.

For problem 5, if you can re-create the problem, start the buffer trace with the `ID=ISTTOPAG` operand and re-create the problem.

If you cannot distinguish between problem 4 on page 330 and 5 on page 330 and the VTAM topology agent abended, try to re-create the problem and start the buffer trace with the ID=ISTTOPAG and ID=ISTNOTIF operands.

If you do not know what caused the problem, start the buffer trace with all three values for the ID operand.

In Figure 40, you can see what is traced when each value for the ID operand is specified for the user buffer trace.

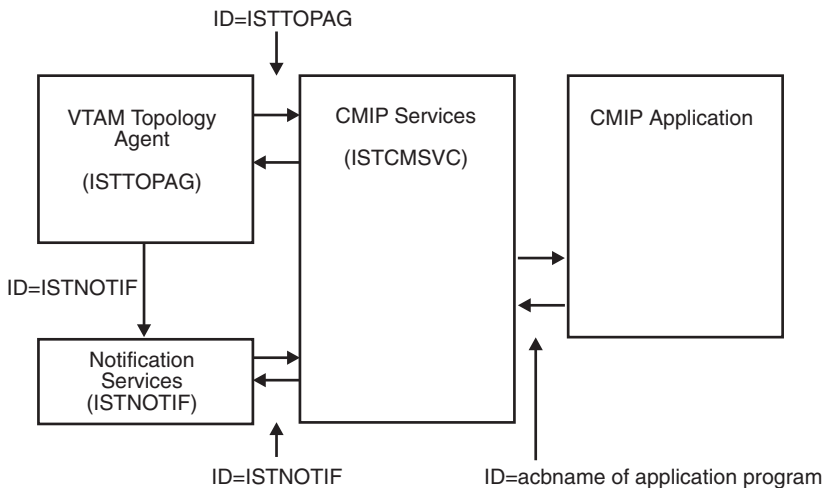


Figure 40. Trace points for the user buffer trace for CMIP services

Choosing between partial and full buffer contents trace

You can request a partial buffer contents trace or a full buffer contents trace. For a partial buffer contents trace, VTAM writes trace records with a maximum size of 256 bytes. The partial buffer contents trace is the default.

Because of the possible effect on storage and performance, use the full buffer trace only when you need complete buffer data for problem determination or when you are tracing the CMIP services API. The full buffer contents trace can increase storage use due to the larger size of trace records and the possible need to increase the size of the trace data set. Performance may be degraded due to the additional time needed to write the complete buffer trace records.

For a full buffer contents trace, VTAM records all of the data transmitted in message buffers. The full buffer trace record has a maximum length of 8K bytes including trace field headers, transmission headers, request/response headers, and data. If storage is not available to record a trace record in an 8K block, VTAM will record the trace record in 256-byte blocks until the complete trace record is recorded.

Buffer contents trace operation

Make sure that the GTF with the TRACE=USR option is active before starting this trace. See "Starting the generalized trace facility (GTF)" on page 322.

If you are using VTAM CMIP services or the VTAM topology agent, a large amount of data is traced. Therefore, you should increase the value of the GTF BLOK parameter and the GTF REGION parameter. For example, specify BLOK=400 and REGION=0 (to indicate no limit on region size). You should also specify the AMOUNT=FULL operand on the TRACE command for the buffer trace so that the entire CMIP string is traced.

Note: If the GTF does not have enough storage to record all the segments, it might record only the last segment. The GTF will report the number of lost records. You might want to increase the storage available to the GTF to avoid losing records.

Start the buffer contents trace with the MODIFY TRACE command or the TRACE start option.

Specify:	To trace:
ID= <i>nodename</i>	Requests and responses flowing between VTAM and <i>nodename</i> .
ID=ISTNOTIF	Event notification activity between the VTAM topology agent and notification services of CMIP services. Specify the AMOUNT=FULL operand when you start the trace.
ID=ISTPUS	Request and response units (RUs) for sessions between the host physical unit and another physical unit type 4 or 5. (These RUs include ER-ACT, ER-OP, and ER-TEST RUs.)
ID=ISTTOPAG	Requests and responses to and from the VTAM topology agent. Specify the AMOUNT=FULL operand when you start the trace.
ID=VTAM	Request and response units (RUs) for SSCP sessions.

If you use the SCOPE=ALL or EVERY operand when you start the trace, VTAM records messages to and from the specified node and all its valid subnodes.

Note: You cannot use the SCOPE=ALL or EVERY operands with ID=ISTPUS.

For more information on the MODIFY TRACE command, see *z/OS Communications Server: SNA Operation*. For more information on the TRACE start option, see *z/OS Communications Server: SNA Resource Definition Reference*.

To format and print buffer contents trace output, use ACF/TAP or IPCS.

For more information, see "Formatting and printing trace records" on page 323.

Buffer contents trace output

Figure 41 on page 333 shows an example of partial buffer contents trace output. Figure 42 on page 333 shows an example of full buffer contents trace output. Figure 43 on page 334 shows an example of full buffer contents trace output for the CMIP services API. The example shows the first few lines of a CMIP string in the buffer trace. ISTCMSVC is VTAM CMIP services. This application program name appears in the buffer trace as the origin application program name in this example. The invoke ID field is X'0004000B' in this example.


```

BUFF APPL12 /TERM200 LRC(000,000) INBOUND
VTAM TH=40000000 200000C2 0000 000C 00000004 IC000005 0028004 0040 RH=030080
      7DD8E811 D7F1F3C8 C1E3C140 C9E240C3 *'QY.PITHIS DATA IS C*
      D6D4D4C9 D5C740C9 D540D6D5 40E4E2C5 *OMING IN ON THE USE *

BUFF TERM200 /APPL12 LRC(000,000) OUTBOUND
USER C5D5E3D9 40C4C1E3 C140E3D6 40C5C8D6 *ENTER DATA TO ECHO B*
      C5D3D6E6 4B5CF04B D7D9C5E2 E240C5D5 *ELOW.*0.PRESS ENTER *

```

Figure 41. Example of partial buffer contents trace output

```

USRFD FEF ASCB 00EBD380          JOBN ECH042A
      BUFF  NETA.APPL1          /NETA.APPL2          LRC(000,000)  INBOUND  COMPLETE SEGMENT
      VTAM  TH=40000000 00000000 00000001 00000001 1C000014 001B08A7 0003  RH=838000
            GMT-11/02/93 17:32:48.795383  LOC-11/02/93 13:32:48.701175

USRFD FEF ASCB 00EBD380          JOBN ECH042A
      BUFF  NETA.APPL1          /NETA.APPL2          LRC(000,000)  INBOUND  FIRST  SEGMENT  SEQ(001)
      VTAM  TH=40000000 00000000 00000001 00000001 1C000014 001B08A7 232B  RH=0380C0
            C1D7D7D3 F1404040 0000D5C5 E3C14040 40400000 *APPL1  ..NETA  ..*
            40D9C5C1 C4E840C6 D6D940C6 C9D9E2E3 40C9D5D7 * READY FOR FIRST INP*
            E4E34B40 D3D6C7D6 D540C4C1 E3C1407E 40E2C9D4 *UT. LOGON DATA = SIM*
            D3D6C7D6 D5150000 00000000 00000000 00000000 *LOGON.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            :
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 *.....*
            GMT-11/02/93 17:32:48.802207  LOC-11/02/93 13:32:48.707999

USRFD FEF ASCB 00EBD380          JOBN ECH042A
      BUFF  NETA.APPL1          /NETA.APPL2          LRC(000,000)  INBOUND  LAST  SEGMENT  SEQ(002)
      VTAM  00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            :
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00000000 00000000 00000000 *.....*
            00000000 00000000 00 *.....*
            GMT-11/02/93 17:32:48.802423  LOC-11/02/93 13:32:48.708215

```

Figure 42. Example of full buffer contents trace output

```

BUFF USER      NETA.ISTTOPAG  /NETA.ISTCMSVC      LRC(000,000)    OUTBOUND    COMPLETE SEGMENT
00000100 0004000B 00000002 00000001 2EB911F5 *.....5*
00000000 00000000 07537408 A2998360 A3A89785 *.....src-type*
40F16B40 A2998340 81F16B40 8485A2A3 60A3A897 * 1, src al, dest-typ*
8540F06B 8485A2A3 407DF14B F34BF1F8 4BF04BF2 *e 0,dest '1.3.18.0.2*
4BF44BF6 7ED5C5E3 C15EF14B F34BF1F8 4BF04BF0 *.4.6=NETA;1.3.18.0.0*
4BF2F0F3 F27EC1F0 F2D55EF1 4BF34BF1 F84BF04B *.2032=A02N;1.3.18.0.*
F04BF2F2 F1F67E4D A2A39989 9587407F E29581D3 *0.2216=(string "Snal*
96838193 E3969796 939687A8 7F5D7D6B 94A28740 *ocalTopology")',msg *
C3D4C9D7 60F14BD9 D6C9E581 9784A440 4D8995A5 *CMIP-1.R0IVapdu (inv*
969285C9 C440F2F6 F2F1F5F5 6B409697 859981A3 *okeID 262155, operat*
89969560 A58193A4 8540F76B 40819987 A4948595 *ion-value 7, argumen*
A3404D82 81A285D4 81958187 8584D682 918583A3 *t (baseManagedObject*
C39381A2 A240F14B F34BF1F8 4BF04BF0 4BF2F1F5 *Class 1.3.18.0.0.215*
F26B4082 81A285D4 81958187 8584D682 918583A3 *2, baseManagedObject*
C995A2A3 81958385 404D8489 A2A38995 87A489A2 *Instance (distinguis*
888584D5 81948540 4DD98593 81A389A5 85C489A2 *hedName (RelativeDis*
A3899587 A489A288 8584D581 9485404D C1A3A399 *tinguishedName (Attr*
8982A4A3 85E58193 A485C1A2 A28599A3 89969540 *ibuteValueAssertion *
4D81A3A3 998982A4 A385E3A8 978540F1 4BF34BF1 *(attributeType 1.3.1*
F84BF04B F24BF44B F66B4081 A3A39989 82A4A385 *8.0.2.4.6, attribute*
E58193A4 85407FD5 C5E3C17F 5D5D6B40 D9859381 *Value "NETA")), Rela*
A389A585 C489A2A3 899587A4 89A28885 84D58194 *tiveDistinguishedNam*
85404DC1 A3A39989 82A4A385 E58193A4 85C1A2A2 *e (AttributeValueAss*
8599A389 9695404D 81A3A399 8982A4A3 85E3A897 *ertion (attributeTyp*

```

Figure 43. Example of full buffer contents trace output for CMIP services API

Fields in the buffer contents trace

Table 17 explains the trace fields. In addition to these fields, operating-system-dependent fields may appear. For a description of these fields, see Table 16 on page 325.

Table 17. Fields in the buffer contents trace

Field header	Meaning
BUFF <i>destname/origname</i>	The destination (<i>destname</i>) and origin (<i>origname</i>) node name. In Figure 41 on page 333, the destination is APPL12 and the origin is TERM200.
LRC(<i>xxx,yyy</i>)	The number of records lost since the last trace record was written because the trace facility could not get a VTAM buffer. <i>xxx</i> is the destination's lost record count, and <i>yyy</i> is the source's lost record count.
INBOUND or OUTBOUND	The direction of the traced data (inbound or outbound) with respect to this host subarea. By use of this field and the PIU sequence number in the TH, requests and corresponding responses can be matched.
<i>position</i> SEGMENT (full buffer trace only)	Indicates whether this trace record is FIRST, MIDDLE, or LAST in a series of trace records generated for one trace invocation. If only one trace record is needed, the value is COMPLETE. The segment indicator field appears only when full buffer contents tracing is in effect.
SEQ(<i>xxx</i>) (full buffer trace only)	A sequence number indicating the sequence in which trace records were generated. The sequence number appears only when a series of trace records is generated for a single trace invocation. The sequence number does not appear when one trace record shows a complete buffer. The sequence number starts at 1, and upon reaching 255, wraps to 0. A gap in sequence numbers could indicate lost trace records. The sequence number field applies only when full buffer contents tracing is in effect.
VTAM or USER	Indicates where the message buffers were traced. VTAM means that the buffers were traced in TSC (in which case the TH and the RH are included in the trace record). USER means that the buffers were traced in API (in which case the TH and the RH are not included).
TH	The transmission header portion of the path information unit (PIU).

Table 17. Fields in the buffer contents trace (continued)

Field header	Meaning
RH	The request/response header portion of the PIU.

Note: The rest of the trace record shows the contents of the buffer as displayed in Figure 41 on page 333.

Notes:

1. If the buffer trace information is out of sequence, your trace could have wrapped.
2. User entries are not printed by ACF/TAP.
3. Confidential data is *not* recorded in trace records. When the trace facility detects confidential data (CONFTXT=YES is specified on the application program's NIB macroinstruction), the user data is replaced with the marker in the trace output.

The marker is **CONFIDENTIAL AND SUPPRESSED**.

If you are using the VTAM encryption facility, data on a cryptographic session is handled in the same way as confidential data.

Configuration services XID exit (CSX) buffer trace

The configuration services XID exit (CSX) buffer trace shows the input and output of the CSX ISTECCS.

Trace points are invoked before and after CSX execution, and the following exit call functions can be traced:

- Begin
- XIDs for DYNAMIC PUs
- XIDs for PREDEFINED PUs
- Connection Status
- Failure
- End

See *z/OS Communications Server: SNA Customization* for more information on configuration services exit.

When to use the configuration services XID exit (CSX) buffer trace

To avoid impacts to performance from the amount of data generated by this trace, use the CSX buffer trace to diagnose suspected problems with your CSX code.

Configuration services XID exit (CSX) buffer trace operation

Start the CSX buffer trace with the MODIFY TRACE command. For more information about the MODIFY TRACE command, see *z/OS Communications Server: SNA Operation*.

Make sure that the GTF with the TRACE=USR or TRACE=USRP option is active before starting this trace. See "Starting the generalized trace facility (GTF)" on page 322. To format and print the data recorded by the GTF, use IPCS and set USR(FF1) on the GTFTRACE option.

For more information on printing trace output, see "Formatting and printing trace records" on page 323.

Configuration services XID exit (CSX) trace record output

Output is formatted into three or four sections:

- The first section is the register save area (ISTRSA). It is formatted using displacements instead of its virtual storage address.
- The second section is the parmlist. It is also formatted using displacements instead of the virtual storage address.
- The third section is the area of storage that contains most of the input data to the exit, which is used by most of the CSX functions.
- The fourth section is optional, and can contain the PARMS=<string> that was entered on the MODIFY EXIT operator command.

Figure 44 shows an example of CSX buffer trace output.

```

EXIT          TO: ISTECCS      FROM: SSCP1A      TYPE: EXITCALL
SAVEAREA
00000000    C9C5C3E2 03C38C58 03C702A8 844C16F0 *IECS.C...G.yd<.0*
00000010    00000088 00000000 00025CA0 03BEE858 *...h.....*...Y.*
00000020    03C38810 83BEE528 00000080 03C702A8 *.Ch.c.V.....G.y*
00000030    0000007F FFFFFFFF 844C0E80 00A6F0F8 *..."..."d<...w08*
00000040    044C1E29 844C0E2A                *.<..d<..      *

PARMLIST
00000000    00025CB8 00025CBA 00025CBB 00025C98 *..*...*...*...*q*
00000010    007D9290                *. 'k.

STORAGE
00025CB8    00160000 00D5C5E3 C14BE2E2 C3D7F1C1 *.....NETA.SSCP1A*
00025CC8    40404040 4040                *                *

PARMS=
007D9290    0013C689 99A2A340 8183A389 A581A389 *..First activati*
007D92A0    96955A                *on.

EXIT          TO: SSCP1A          FROM: ISTECCS      TYPE: EXITRETN
SAVEAREA
00000000    C9C5C3E2 03C38C58 03C702A8 844C1A2C *IECS.C...G.yd<..*
00000010    00000000 00000000 00A65C54 03BEE858 *.....w*...Y.*
00000020    00A65C54 83BEE528 00000080 03C702A8 *.w*.c.V.....G.y*
00000030    0000007F FFFFFFFF 844C0E00 00A6F0F8 *..."..."d<...w08*
00000040    044C1E29 844C0E2A                *.<..d<..      *

PARMLIST
00000000    00025CB8                *..*..          *

STORAGE
00025CB8    00160080 00D5C5E3 C14BE2E2 C3D7F1C1 *.....NETA.SSCP1A*
00025CC8    40404040 4040                *                *

```

Figure 44. Example of configuration services XID exit (CSX) buffer trace output

Directory services session management exit (DSME) buffer trace

The directory services session management exit (DSME) buffer trace shows the input and output of the DSME ISTECCDM.

Trace points are invoked before and after DSME execution, and the following exit call functions can be traced:

- Begin
- Border Node Selection
- Central Directory Server Selection
- Central Resource Registration Selection
- Interchange Node Selection
- End
- Initial Authorization
- Exit Replacement

See *z/OS Communications Server: SNA Customization* for more information on directory services management exit routines.

When to use the directory services session management exit (DSME) buffer trace

To avoid impacts to performance from the amount of data generated by this trace, use the DSME buffer trace to diagnose suspected problems with your DSME code.

Directory services session management exit (DSME) buffer trace operation

Start the DSME buffer trace with the `MODIFY TRACE` command. For more information about the `MODIFY TRACE` command, see *z/OS Communications Server: SNA Operation*.

Make sure that the GTF with the `TRACE=USR` or `TRACE=USRP` option is active before starting this trace. See “Starting the generalized trace facility (GTF)” on page 322. To format and print the data recorded by the GTF, use `IPCS` and set `USR(FF1)` on the `GTFTRACE` option.

For more information on printing trace output, see “Formatting and printing trace records” on page 323.

Directory services session management exit formatted output

Output is formatted into several sections:

- The first section is the register save area (ISTRSA). It is formatted using displacements instead of its virtual storage address.
- The second section is the parmlist. It is also formatted using displacements instead of the virtual storage address.
- The next section is the 8-byte user field.
- The next section is the environment vectors.
- The next section is the function code and related search information.
- The next section is the user data section.
- The next section contains either the exit options (for `BEGIN`), the OLU information structure, the CDS list (for `CRR` selection), or the `PARMS= <string>` (if `PARMS=` was specified on the `VTAM MODIFY EXIT` command for `OPT=REPL` or `END`).

Note: For functions `REPL`, `END`, and `CRR`, this would be the end of the trace record.

- The next section would contain the border node options (for BEGIN), or the DLU information structure.
- The next section contains either the PARMs= <string> (for BEGIN, if PARMs= was specified on the VTAM MODIFY EXIT command for OPT=ACT), or the network qualified CP name vector.

Note: For the BEGIN function, this would be the end of the trace record.

- The next section would contain the search correlator structure.
- The next section would contain the PCID modifier structure.

Note: For the INITIAL AUTHORIZATION function, this would be the end of the trace record.

- The last section would contain either the subnetwork routing list (for BN selection), interchange node list (for ICN selection), or the CDS list (for CDS or ADS selection).

Figure 45 on page 339 shows an example of DSME buffer trace output.

```

EXIT          TO: ISTECDM      FROM: SSCP1A      TYPE: EXITCALL
SAVEAREA
00000000 C9C5C4D4 03C3DC58 03B19438 8450CD6C *IEDM.C....m.d&.*
00000010 83A334D8 00000008 03C3DF80 00A700F8 *ct.Q.....C...x.8*
00000020 03C3D810 0427AF58 00000080 00000080 *.CQ.....*
00000030 0000007F FFFFFFFF7F 8450C480 0450D429 *..."..."d&D&M.*
00000040 0450CD6C 8450C42A *.&.%d&
D          *

PARMLIST
00000000 03BF2C88 00000000 00000000 03C3DFA4 *...h.....C.u*
00000010 03C3DFC0 03C3DFC6 03C3DFCA 03C3DFCC *.C...C.F.C...C...*
00000020 807D9290 *. 'k. *

USER_FLD
03BF2C88 00000000 00000000 *. . . . . *

ENV_VECT
03C3DFA4 001C0606 D5C5E3C1 0807E2E2 C3D7F1C1 *...NETA..SSCP1A*
03C3DFB4 040CBC00 00000000 00000000 *. . . . . *

FNCTCODE
03C3DFC0 FE404040 4040 *. *

USERDATA
03C3DFC6 00000000 *. . . . . *

EXITOPTS
03C3DFCA 00000080 *. . . . . *

PARMS=
807D9290 001C0606 E3C5E7C3 C4D440C1 83A389A5 *..ISTECDM Activ*
807D92A0 81A38996 954B4B4B *.ation... *

EXIT          TO: SSCP1A      FROM: ISTECDM      TYPE: EXITRETN
00000000 C9C5C4D4 03C3DC58 03B19438 8450CE8E *IEDM.C....m.d&.*
00000010 00000000 00000008 03C3DF80 00A700F8 *.....C...x.8*
00000020 03C3DF80 0427AF58 00000080 00000080 *.C.....*
00000030 0000007F FFFFFFFF7F 8450C400 0450D429 *..."..."d&D&M*
00000040 0450CE8E 8450C42A *.&..d&
D          *

PARMLIST
00000000 03BF2C88 00000000 00000000 03C3DFA4 *...h.....C.u*
00000010 03C3DFC0 03C3DFC6 03C3DFCA 03C3DFCC *.C...C.F.C...C...*
00000020 807D9290 *. 'k. *

USER_FLD
03BF2C88 01010101 00000001 *. . . . . *

ENV_VECT
03C3DFA4 001C0606 D5C5E3C1 0807E2E2 C3D7F1C1 *...NETA..SSCP1A*
03C3DFB4 040CBC00 00000000 00000000 *. . . . . *

FNCTCODE
03C3DFC0 FE404040 4040 *. *

USERDATA
03C3DFC6 00000000 *. . . . . *

EXITOPTS
03C3DFCA 9F020180 *. . . . . *

```

Figure 45. Example of directory services session management exit (DSME) buffer trace output

I/O trace

The I/O trace shows requests and responses that flow between VTAM and network nodes. You can trace I/O activity for any of the following types of nodes:

- Application program
- Physical unit
- Logical unit
- SNA cluster controller
- NCP
- SSCP
- Host physical unit
- Host as an intermediate routing node
- Channel attachment major node
- Cross-domain resource
- Cross-domain resource manager
- RTP pipe
- TRLE

Restriction: I/O trace is not supported for a TRLE that represents an IBM 10GbE RoCE Express interface.

The maximum I/O trace record length is 272 bytes.

Note:

1. If you want to trace a session between an LU and an application program, you must start the trace at the host where the application program resides.
2. I/O trace records are not recorded for conversation level data exchanged between two VTAM/APPC applications residing on the same host and using the APPCCMD macroinstruction interface to communicate.
3. I/O trace provides packet tracing capability for OSA-Express QDIO and HiperSockets data devices because CCW trace does not exist for these devices. Packet trace for OSA-Express QDIO and HiperSockets will appear as ODPK records in the external VIT. A length field is provided on the MODIFY TRACE command for OSA-Express QDIO and HiperSockets devices to override the existing 272-byte trace limit for I/O trace.
4. Do not enable I/O trace for an OSA-Express2 or later data device that is used to capture OSA-Express network traffic analyzer trace data. The VARY TCPIP,OSAENTA command described in z/OS Communications Server: IP Diagnosis Guide has its own ability to filter, capture, and format this data. If I/O trace is enabled for a data device used for capturing trace data, only the first 28 bytes of each packet are traced.
5. You must use a combination of the TCP/IP packet trace facility and VTAM internal trace (VIT) records to analyze Shared Memory Communications - RDMA (SMC-R) link traffic. The RPST records in the VIT represent data being sent outbound by using SMC-R communications. The RPLR records in the VIT represent data arriving inbound by using SMC-R communications. For information about the TCP/IP packet trace, see z/OS Communications Server: IP Programmer's Guide and Reference.

When to use the I/O trace

Use the I/O trace to record the order that PIUs flow between network nodes and VTAM. For example, you might use this trace to determine whether an application

program receives all the responses that it should and whether VTAM forwards all the requests issued by the application program.

The I/O trace is now done using the external VIT. Data items are formatted like VIT external trace entries.

I/O trace operation

Before starting the I/O trace, make sure that the GTF with the TRACE=USR option is active. See “Starting the generalized trace facility (GTF)” on page 322.

Start the I/O trace with the MODIFY TRACE command or the TRACE start option.

Specify:	To trace:
ID= <i>nodename</i>	Requests and responses flowing between VTAM and <i>nodename</i> .
ID=VTAM	Request and response units (RUs) for SSCP sessions.
ID=ISTPUS	Request and response units (RUs) for sessions between the host physical unit and another physical unit type 4 or 5 (these RUs include ER-ACT, ER-OP, and ER-TEST RUs).
ID=ISTIRN	Request and response units (RUs) that flow through this host while this host is acting as an intermediate routing node.

Notes:

1. If you use the SCOPE=ALL or EVERY operand when you start the trace, the trace contains I/O activity for the specified node and all its valid subnodes. You *must* specify SCOPE=ALL when tracing a channel-attachment major node or when tracing an APPN PU. You *cannot* use the SCOPE=ALL or EVERY operands with ID=ISTPUS or ID=ISTIRN.
 2. You may trace a link in a channel-attachment major node but not a link station.
-

See z/OS Communications Server: SNA Operation for information on the MODIFY command. See z/OS Communications Server: SNA Resource Definition Reference for more information on the TRACE start option.

QDIOSYNC trace

The QDIOSYNC trace is not a traditional trace in which output is generated based on specific events. Instead, the QDIOSYNC trace freezes and captures (logs) OSA-Express2 or later diagnostic data in a timely manner. In addition to or instead of using the hardware management console (HMC) to manually capture the diagnostic data, you can arm the OSA-Express2 or later adapter to automatically capture diagnostic data when one of the following conditions occurs:

- The adapter detects an unexpected loss of host connectivity. Unexpected loss of host connectivity occurs when the adapter receives an unexpected halt signal from the host or when the host is unresponsive to OSA requests. An unexpected halt signal includes VTAM InOp traps (for example, ISTLLCIE InOpCode 101).
- The adapter receives a CAPTURE signal from the host. A CAPTURE signal is sent by the host when one of the following conditions occurs:
 - The VTAM-supplied message processing facility (MPF) exit (IUTLLCMP) is driven.
 - Either the VTAM or TCP/IP functional recovery routine (FRR) is driven with ABEND06F. ABEND06F is the result of a SLIP PER trap that specifies ACTION=RECOVERY.

Restriction: The SLIP must be a SLIP PER trap to specify ACTION=RECOVERY.

When arming an OSA-Express2 or later adapter for QDIOSYNC, you can specify an optional filter that alters what type of diagnostic data that the adapter collects. This filtering reduces the overall amount of diagnostic data collected and therefore decreases the likelihood that pertinent data is lost.

Arming an OSA-Express2 or later adapter has no effect on the host or OSA-Express2 or later performance. However, using a PER-type SLIP trap generally increases host CPU utilization.

When to use the QDIOSYNC trace

Use the QDIOSYNC trace when you are re-creating a problem with an OSA-Express2 or later connection. QDIOSYNC enables automatic capture of OSA-Express2 or later diagnostic data without using the HMC console, reducing the likelihood of overwriting pertinent diagnostic data. QDIOSYNC also uses other system facilities to enable host-initiated capture of OSA-Express2 or later diagnostic data and to initiate a host dump simultaneously with the host-initiated capture of OSA-Express2 or later diagnostic data. These other system facilities include the z/OS MPF exit facility and the SLIP PER trap command.

Specifying the VTAM-supplied MPF exit routine module name [USEREXIT(IUTLLCMP)] in the MPFLSTxx member of SYS1.PARMLIB and activating that member result in a CAPTURE signal being sent to all armed OSA-Express2 or later adapters when the corresponding message is issued. See z/OS MVS Installation Exits for more information about the use of MPF and the MPF PARMLIB member.

Specifying ACTION=RECOVERY on a SLIP PER trap command drives the executing FRR with an ABEND06F abend when the SLIP trap is triggered. Both the VTAM and TCP/IP FRRs detect the ABEND06F abend and initiate sending a CAPTURE signal to all armed OSA-Express2 or later adapters.

Do not use QDIOSYNC to unconditionally arm an OSA-Express2 or later adapter when it is shared by other operating systems and those operating systems might use this function. In this case, the function should be coordinated between all sharing operating systems.

QDIOSYNC trace operation

The VTAM MODIFY TRACE and NOTRACE start options and commands, with the value TYPE=QDIOSYNC, are used to activate and terminate QDIOSYNC trace. As with most other TRACE commands, you can save a QDIOSYNC command so that if the TRLE defining the OSA-Express2 or later devices to be synchronized is not active, the command is applied when the TRLE is activated. Unlike other TRACE and NOTRACE commands, the value ID=* is supported. When the value ID=* is specified with SAVE=NO, the ID=* value indicates that the QDIOSYNC command is to be applied to all currently active TRLEs that define OSA-Express2 or later adapters. When the value ID=* is specified with SAVE=YES, the ID=* value indicates that the QDIOSYNC command is to be applied to all currently active TRLEs that define OSA-Express2 or later adapters and to those that are activated by this VTAM in the future.

See MODIFY NOTRACE and MODIFY TRACE commands in z/OS Communications Server: SNA Operation and TRACE start option in z/OS Communications Server: SNA Resource Definition Reference for additional details about employing QDIOSYNC using TRACE.

You can use the VTAM DISPLAY TRACES, DISPLAY TRL, and DISPLAY ID commands to determine the current QDIOSYNC usage. See DISPLAY TRACES, DISPLAY TRL, and DISPLAY ID commands in z/OS Communications Server: SNA Operation for additional details.

Using QDIOSYNC to synchronize OSA diagnostic data with host diagnostic data:

You can use QDIOSYNC to synchronize OSA diagnostic data with host diagnostic data.

Procedure

Perform the following steps to use QDIOSYNC to synchronize OSA diagnostic data with host diagnostic data:

1. Determine which OSA-Express2 or later adapters must be armed. You can use the TRACE command or TRACE start option to arm the adapters. The advantage to using the command is that VTAM does not need to be restarted for the command to take effect. The advantage to using the start option is that you must code it only once.
2. Determine what value you will use for the OPTION operand on the TRACE command.

Use the OPTION value to filter the OSA trace table records. The OPTION value specifies which data devices (controlled by the adapters that you are arming) will have their activity traced and in what direction. Possible values are:

ALLIN

Directs OSA to collect only inbound diagnostic data for all devices.

ALLOUT

Directs OSA to collect only outbound diagnostic data for all devices.

ALLINOUT

Directs OSA to collect inbound and outbound diagnostic data for all devices.

Tip: Specifying the ALLIN, ALLOUT, or ALLINOUT filters dictates recording of activity from data devices not only controlled by this z/OS Communications Server but also by other operating systems sharing the OSA-Express2 or later adapter.

IN Directs OSA to collect only inbound diagnostic data for devices defined to this VTAM.

OUT Directs OSA to collect only outbound diagnostic data for devices defined to this VTAM.

INOUT

Directs OSA to collect inbound and outbound diagnostic for devices defined to this VTAM.

Tip: Specifying the IN, OUT, or INOUT filters dictates recording of activity from data devices controlled by this z/OS Communications Server only. Activity from data devices controlled by other operating systems sharing the OSA-Express2 or later adapter is not recorded.

If you cannot determine the scope and filtering to use, use the default of ALLINOUT.

3. Issue the TRACE command to arm one or more adapters, or to arm all adapters. Take one of the following steps:

- To individually arm the adapters, issue the TRACE ID=*trle_name* OPTION=*option* command for each adapter that you want to arm.
 - To arm all adapters using a single command, issue the TRACE ID=* OPTION=*option* command. If you have several OSAs to arm but you do not want to arm all of them, it might be easier to arm all OSAs and then individually disarm those you do not want armed.
4. Issue the DISPLAY TRACES command and optionally the DISPLAY TRL or DISPLAY ID command to check their QDIOSYNC status.
 5. For message ID traps, activate the MPF parmlib member, which specifies USEREXIT(IUTLLCMP).

* This MPFLSTxx identifies the messages which lead to capture of
 * armed OSA-Express devices. If any of the following message are
 * issued, IUTLLCMP (VTAM provided MPF exit) gains control and
 * schedules the capture of all armed OSA-Express devices.
 *

```
EZZ4343I ERROR xxxx REGISTERING IP ADDRESS<IP Addr> FOR ...
EZZ4339I INTERFACE interface_name FAILED - ADAPTER SIGNAL ...
EZZ4327I ERROR xxxx REGISTERING IP ADDRESS
EZZ4328I ERROR xxxx SETTING ROUTING FOR DEVICE
EZZ4343I,SUP(NO),USEREXIT(IUTLLCMP)
EZZ4339I,SUP(NO),USEREXIT(IUTLLCMP)
EZZ4327I,SUP(NO),USEREXIT(IUTLLCMP)
EZZ4328I,SUP(NO),USEREXIT(IUTLLCMP)
```

- a. To activate this parmlib member, issue SET MPF=(*xx,zz*) (where *xx* is a new parmlib member and *zz* is your current MPFLST*zz* parmlib member).
 This captures all armed OSA-Express2 or later devices when any of the four TCP/IP messages (EZZ4343I, EZZ4339I, EZZ4327I, or EZZ4328I) are issued.
- b. You should also set a corresponding SLIP trap for each message in the parmlib member to initiate a host dump.

Example:

```
SL DEL, ID=MEZ1, END
SL SET, ID=MEZ1, MSGID=EZZ4343I, A=(STOPGTF, SVCD), MATCHLIM=1,
JOBLIST=(TCP*, NET*),
DSPNAME=('TCP*'.*, 01.CSM*, 'NET*'.IST*),
SDATA=(RGN, ALLNUC, CSA, LSQA, PSA, SQA, SUM, SWA, TRT, LPA),
END
```

This SLIP will be triggered when the EZZ4343I TCP/IP error message is issued. Because the MPF exit is active for EZZ4343I, all armed OSA-Express2 or later adapters will be sent a CAPTURE signal. Each armed adapter freezes the diagnostic data and writes it to the hard disk of the HMC. An SVC dump will also be taken with the title SLIP DUMP=MEZ1. This is a sample. If you choose to use it you must ensure that your job and data space correspond to the parameters in the sample.

Tips:

- The first command in this example deletes any SLIP trap with the name MEZ1.
- The *jobname* value for VTAM is NET and the *jobname* for TCP/IP is TCP (Use your own job name values).
- All data spaces created by TCP/IP are dumped.
- All data spaces created by the master scheduler that contain CSM in the name are dumped.
- All data spaces created by VTAM whose names start with IST are dumped.
- GTF is also stopped if it was running when this SLIP trap is matched.

- For module offset trap problems, you must code a SLIP PER trap on a specific module. The following example assumes that the abend occurred in module ISTLLCIE. You first have to find the starting address of ISTLLCIE by issuing the following command:

```
D NET,VTAMSTOR,MOD=ISTLLCIE

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = VTAMSTOR 992
IST1571I ISTLLCIE ENTRY POINT IS 3A412CBC LEVEL IS HVT6180
IST1574I -020 47F0F01C 17C9E2E3 D3D3C3C9 * .00..ISTLLCI
IST1574I -010 C540F0F5 4BF2F3F1 40C8E5E3 F6F1F8F0 *E 05.231 HVT6180
IST1574I +000 90ECD00C 05C041B0 CFFF4840 02041884 *..... ..D
IST1574I +010 8B400007 1F4858A0 04088B40 00065E40 *. .... ..;
IST314I END
```

In this example, assuming that you want to gather documentation at +200'x' in ISTLLCIE, code the SLIP as follows:

```
SL DEL, ID=MEZ2, END
SL SET, IF, ID=MEZ2, RA=(3A412EBC), A=(STOPGTF, RECOVERY, SVCD),
MATCHLIM=1, JOBLIST=(TCP*, NET*),
DSPNAME=('TCP*', *, 01.CSM*, 'NET*'.IST*),
SDATA=(RGN, ALLNUC, CSA, LSQA, PSA, SQA, SUM, SWA, TRT, LPA),
END
```

Guideline: The MEZ2 trap is an example; you must locate a different module and determine the instruction address by adding a different offset. The example is provided not only to show how to determine the instruction address but also to stress the use of the RECOVERY parameter in the ACTION list. If you choose to use this example you must ensure that your job and data space names correspond to the parameters in the sample.

Tips:

- The first command in this example deletes any SLIP trap with the name MEZ2.
 - The SLIP PARM RA=(3A412EBC) is 200'x' bytes into ISTLLCIE. In this example, ISTLLCIE is loaded in common storage, so additional SLIP parameters are not needed (for example, the *jobname* parameter).
 - The *jobname* value for VTAM is NET and the *jobname* value for TCP/IP is TCP (Use your own job name values.)
 - All data spaces created by TCP/IP are dumped.
 - All data spaces created by the master scheduler that contain CSM in the name are dumped.
 - All data spaces created by VTAM whose names start with IST are dumped.
 - GTF is also stopped if it was running when this SLIP trap is matched.
- Re-create the problem.
 - Optionally, disarm any or all adapters using the MODIFY NOTRACE command.
 - Use the HMC to locate and copy the OSA-Express2 or later diagnostic data.

Resource state trace

The resource state trace creates VTAM internal traces (VIT) entries when the current state or desired state, or both, of a resource for which tracing has been requested changes. You can choose to trace all resources, specific resources, or all resources of a particular type, for example, all the application programs.

When to use the resource state trace

To avoid impacts to performance from the additional VIT entries that are generated, use resource state tracing for diagnosing specific problems and try to narrow the number of resources being traced.

Resource state trace operation

Start the resource state trace with the MODIFY TRACE command, or use the TRACE start option with TYPE=STATE specified.

Specify:	To Trace:
ID= <i>resourcename</i>	Changes to the current state or desired state, or both, of <i>resourcename</i> .
OPTION= <i>options</i>	Changes to the current state or desired state, or both, of all resources of type <i>option</i> .

Note: If you use the OPTION=ALL operand, the states of all resource types in your network will be traced.

For a description of the VIT entries created by the resource state trace, see the CSx VIT entry in *z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps* and the VIT. For more information about the MODIFY TRACE command, see *z/OS Communications Server: SNA Operation*. For a description of all the resource states, see *z/OS Communications Server: IP and SNA Codes*.

Session management exit (SME) buffer trace

The session management exit (SME) buffer trace shows the input and output of the session management exit (SME) ISTECA.

Trace points are invoked before and after SME execution, and the following exit functions can be traced:

- Begin function
- Adjacent SSCP selection
- Gateway path list
- Initial authorization
- Secondary authorization
- Initial and final accounting
- Exit replace and replaced
- Virtual route selection
- Alias translation
- XRF session switch
- Adjacent link station selection
- End function
- HPR VR Selection
- HPR RTP Pipe Authorization for the OLU node role
- HPR RTP Pipe Authorization for the DLU node role
- HPR RTP Pipe Authorization for the ANR node role

See *z/OS Communications Server: SNA Customization* for more information on session management exit routines.

When to use the session management exit (SME) buffer trace

To avoid impacts to performance from the amount of data generated by this trace, use the SME buffer trace to diagnose suspected problems with your SME code.

Session management exit (SME) buffer trace operation

Start the SME buffer trace with the MODIFY TRACE command. For more information about the MODIFY TRACE command, see z/OS Communications Server: SNA Operation.

Make sure that the GTF with the TRACE=USR or TRACE=USRP option is active before starting this trace. See “Starting the generalized trace facility (GTF)” on page 322. To format and print the data recorded by the GTF, use IPCS and set USR(FF1) on the GTFTRACE option.

For more information on printing trace output, see “Formatting and printing trace records” on page 323.

Session management exit trace record output

Output is formatted into three or four sections:

- The first section is the register save area (ISTRSA). It is formatted using displacements instead of its virtual storage address.
- The second section is the parmlist. It is also formatted using displacements instead of the virtual storage address.
- The third section is the area of storage that contains most of the input data to the exit, which is used by most of the SME functions. It is listed in dump-like format by virtual address.
- The fourth section is the list of optional data used only for the GWPATH, ADJSSCP, ALIAS, or ALS selection exit functions.

Figure 46 on page 348 shows an example of SME buffer trace output.


```

EXIT      TO: ISTECAA      FROM: SSCP1A      TYPE: EXITCALL

SAVEAREA
00000000 00000000 038E2588 0336BDD0 80009C28 *.....h.....*
00000010 83293E30 00000004 00B4663C 00000006 *C.....*
00000020 00000006 00B4C1F8 00B465E8 00B4663C *.....A8...Y...*
00000030 0346A958 0336BDD0 0346A028 00009488 *..z.....mh*
00000040 00008489 8000748A *..di....*

PARMLIST
00000000 00B46690 00B466CC 00B46914 00B466D4 *.....M*
00000010 00B467CC 00B46B4C 00B468CC 808DB6D0 *.....,<.....*
00000020 00000000 00000000 00000000 008DB6F8 *.....8*
00000030 008DB7A0 00000000 80000000 00000000 *.....*
00000040 00000000 00B46AC8 00B46B04 00B46914 *.....H.,....*
00000050 80B46B08 *.,.,.*

STORAGE
00B46690 00320A06 D5C5E3C1 40404040 0A07E2E2 *....NETA ..SS*
00B466A0 C3D7F1C1 40400A08 C9E2E3D7 E4E24040 *CP1A ..ISTPUS *
00B466B0 08090000 00010000 0A0AD5C5 E3C14040 *.....NETA *
00B466C0 40400000 00000000 00000000 06200000 * .....*
00B466D0 00000000 192D0AC0 0008D9C5 C1D3D5C5 *.....REALNE*
00B466E0 E34008D9 C5C1D3D5 C1D4C508 D5C5E3C1 *T .REALNAME.NETA*
00B466F0 40404040 08D5D6D5 C1D4C540 401A001A * .NONAME ...*
00B46700 001A00FF 1902D9C5 C1D3D5C1 D4C51800 *.....REALNAME..*
00B46710 0000C9E2 E3C3C4D9 C4E80700 0000FE01 *..ISTCDRDY.....*
00B46720 00000000 00000000 00000000 00000000 *.....*
:
00B46B30 00000000 00000000 00000000 00000000 *.....*
00B46B40 00000000 00000000 00000000 EAABEEC3 *.....C*
00B46B50 5353C014 D5C5E3C1 4BE2E2C3 D7F1C140 *....NETA.SSCP1A *
00B46B60 40404040 40000000 0328B000 * .....*

ADJSSCPs
808DB6D0 000A0001 C9E2E3C1 D7D5C3D7 *....ISTAPNCP *

```

Figure 46. Example of session management exit (SME) buffer trace output

SMS (buffer use) trace

The storage management services (SMS) trace records contain information on the use and availability of VTAM buffer pools. SMS trace records are written after a predetermined number of requests occur for VTAM buffers. An IBM-supplied threshold causes a trace record to be written after every 1000 requests.

Note:

1. If the DISPLAY BFRUSE command is issued while this trace is running, the fields MAX TOTAL, MAX USED, and TIMES EXP in the output for DISPLAY BFRUSE reflect buffer usage only since the last trace record was written, because the SMS trace resets these fields.
2. The SMS trace is *not* the same thing as the VTAM internal trace with the SMS option specified. The SMS trace *is* similar to the DISPLAY BFRUSE command. The SMS trace displays in trace output the same information that the DISPLAY BFRUSE command displays on the screen.

When to use the SMS trace

Use the SMS trace during VTAM installation to evaluate VTAM use of buffer pools, to help estimate how many buffers VTAM needs for normal operation, and, with dynamic buffering, to limit buffer pool expansions to peak use periods. You can use the SMS trace with tuning statistics.

SMS trace operation

Start the SMS trace with the MODIFY TRACE command, or use the TRACE start option with TYPE=SMS and ID=VTAMBUF specified. For more information about the MODIFY TRACE command, see z/OS Communications Server: SNA Operation. For more information about the TRACE start option, see z/OS Communications Server: SNA Resource Definition Reference.

Make sure that the GTF with the TRACE=USR option is active before starting this trace. See “Starting the generalized trace facility (GTF)” on page 322. To format and print the data recorded by the GTF, use IPCS and set USR(FF0) or USR(CL02) on the GTFTRACE option.

For more information on printing trace output, see “Formatting and printing trace records” on page 323.

SMS trace record output

Figure 47 shows an example of SMS trace output. The trace fields are explained in Table 18 on page 350. In addition to the fields described here, other operating-system-dependent fields may appear. These fields are described in Table 16 on page 325.

VTAM BUFFERS	MAXU	MAXQ	AVNO	TEXP	MBUF	TOTL
I000	00000015	00000000	00000060	00000000	0000006E	0000006E
BS00	00000034	00000000	00000014	00000000	00000048	00000048
LP00	00000007	00000000	0000003D	00000000	00000040	00000040
XD00	00000004	00000000	0000000B	00000000	0000000F	0000000F
LF00	00000002	00000000	00000076	00000000	00000078	00000078
CRPL	0000003A	00000000	00000096	00000000	000000C8	000000C8
VTAM BUFFERS	MAXU	MAXQ	AVNO	TEXP	MBUF	TOTL
SF00	00000004	00000000	0000003C	00000000	00000040	00000040
SP00	00000000	00000000	0000002A	00000000	0000002A	0000002A
AP00	00000000	00000000	00000038	00000000	00000038	00000038
TI00	00000065	00000000	00000078	00000000	000000B4	000000B4
CRA4	00000002	00000000	0000000C	00000000	0000000C	0000000C
CRA8	00000002	00000000	0000000B	00000000	0000000C	0000000C
VTAM CSAUSE	TOTAL	0030015B	MAX 7FFFFFFF	%FREE QUEUE	00000000	

Figure 47. Example of SMS trace output

Two separate records will be printed, and they may be separated by another trace entry.

Table 18. Fields in the SMS trace

Field header	Meaning
Pool ID	The first field in each record. Pool ID identifies the buffer pool. Pool IDs and their corresponding buffer pool names are: AP Application program pageable pool (APBUF) BS Boundary session block pool (BSBUF) CR Copied RPL pool (CRPLBUF) CRA4 Component recovery area CRA8 Component recovery area IO Fixed I/O pool (IOBUF) LF Large fixed pool (LFBUF) LP Large pageable pool (LPBUF) SF Small fixed pool (SFBUF) SP Small pageable pool (SPBUF) TI HPDT Services (TIBUF) UE User exit control block (UECB) (obsolete) XD XID pool (XDBUF)
MAXU	The maximum number of buffers in the pool that are in use at any time since the last trace record is written.
MAXQ	The maximum number of requests for buffers that are queued waiting for storage at any time since the last trace record is written.
AVNO	The number of available buffers (those not in use at the time the trace record is written).
TEXP	The number of times the buffer pool is expanded since the last trace record is written.
MBUF	The maximum number of buffers that are in the pool at any time since the last trace record is written. This includes both used and unused buffers.
TOTL	The total number of buffers that are in the pool at the time this record is written. This includes both used and unused buffers.
TOTAL	The amount of CSA storage in use by VTAM at the time this record is written.
MAX	The largest amount of CSA storage used by VTAM since the last SMS buffer trace.
FREE QUEUE	The amount of CSA storage allocated to VTAM that is waiting to be freed.
TI	Buffer pool supporting HPDT services
CRA4	Component recovery area (4 KB).
CRA8	Component recovery area (8 KB).

TGET/TPUT trace for TSO/VTAM

The TGET/TPUT trace for TSO/VTAM writes a GTF trace record for each inbound and outbound message that uses the TGET/TPUT/TPG interface (SVC 93) between a TSO command processor and the VTIOC component of TSO/VTAM. Outbound

messages are traced before being placed in the VTIOC queue manager output buffer. Inbound messages are traced before the data is sent to the TSO command processor.

Note: The TGET/TPUT trace does not trace address space ID TPUTs.

When to use the TGET/TPUT trace

Use this trace if the failure is restricted to TSO sessions. This trace can help you determine whether TSO/VTAM or your TSO command processor is causing the problem. For example, if outbound data is correct in the TGET/TPUT trace output, but incorrect in the buffer trace output, the problem is probably in TSO/VTAM or VTAM. Use the following tables as guidelines to determine where the error is occurring:

Table 19. Location of TPUT (outbound) error

Direction of data	If TPUT trace data is:	And buffer trace data is:	Then possible error is in:
Outbound	Correct	Incorrect	<ul style="list-style-type: none"> • VTAM • TSO/VTAM • TPUT option • User edit exits
Outbound	Incorrect	Trace not required	TSO or the command processor
Outbound	Correct	Correct	Network

Table 20. Location of TGET (inbound) error

Direction of data	If TGET trace data is:	And buffer trace data is:	Then possible error is in:
Inbound	Incorrect	Correct	<ul style="list-style-type: none"> • VTAM • TSO/VTAM • TGET option • User edit exits
Inbound	Correct	Trace not required	TSO or the command processor
Inbound	Incorrect	Incorrect	Network

TGET/TPUT trace operation

Start the TGET/TPUT trace with the MODIFY TRACE command and specify TYPE=TSO.

The trace output is a record of inbound and outbound messages for the specified TSO user ID.

Make sure that the GTF with the TRACE=USR option is active before starting this trace. See “Starting the generalized trace facility (GTF)” on page 322.

To print these trace records, use IPCS and specify either USR(FE2) or USR(APTH) on the GTFTRACE option. For more information on printing trace output, see "Formatting and printing trace records" on page 323.

TGET/TPUT trace output

The trace records created by the TGET/TPUT trace have a 12-byte GTF header and a 52-byte trace header followed by the data portion of the RU in unformatted hexadecimal. The entire RU is traced, but will span several trace entries if it is longer than 228 bytes.

Figure 48 shows an example of TGET/TPUT trace output. The trace fields are explained after the figure.

```

IPCS PRINT LOG FOR USER USER1
+
0 **** GTFTRACE DISPLAY OPTIONS IN EFFECT ****
  USR=SEL
0 **** GTF DATA COLLECTION OPTIONS IN EFFECT: ****
  All GTRACE events requested
  RNIO events traced
0 **** GTF TRACING ENVIRONMENT ****
  Release: SP4.2.2  FMID: JBB4422  System name: MVS41D25
  CPU Model: 3090  Version: FF  Serial no. 373247
-HEXFORMAT AID FF FID 00 EID EFE2
+0000 00E8E380 E4E2C5D9 F1404040 E3E2D6D6 E4E30300 C9D5C9E3 E4E2C5D9 F1404040 ..T.USER1 TSOOUT..INITUSER1
+0020 0DD5C5E3 C14BC1F5 F0C9F0F7 F2F14040 40400000 C1114040 1D603CC1 50401D60 .NETA.A50I0721 ..A. .-.A&; :-
+0040 3CC26040 1D603C22 7E40C9C9 C9C94040 C2C2C2C2 40404040 D4404040 D43CC3F0 .B- .-.B= IIII BBBB M M.C0
+0060 401D603C C44F40C9 C9404040 40C240C2 40404040 D4D440D4 D43CC540 401D603C .-.D| II B B MM MM.E .-.
+0080 C55F40C9 C9404040 40C240C2 40404040 D440D440 D43CC650 401D603C C66F40C9 E- II B B M M M.F&; .-.F? I
+00A0 C9404040 40C240C2 40404040 D4404040 D43CC760 401D603C C77E40C9 C9C9C940 I B B M M.G- .-.G= IIII
+00C0 40C2C2C2 C2404040 40D44040 40D43CC8 F0401D60 3C4A4040 1D603C4B 50401D60 BBBB M M.H0 .-.+ .-.&; :-
+00E0 3C4BF240 1DE8C9E2 D7C661D7 C4C63C4C 60401D60 3C4DF040 1D603C4F 40401DE8 ..2 .YISPF/PDF.<- .-.(0 .-.| .Y
+0100 3C4FD340 D3898385 95A28584 .|L Licensed
      GMT-09/03/xx 19:29:49.847841 LOC-09/03/xx 15:29:49.753633
0HEXFORMAT AID FF FID 00 EID EFE2
+0000 00E8E380 E4E2C5D9 F1404040 E3E2D6D6 E4E30300 40D4C9C4 E4E2C5D9 F1404040 ..T.USER1 TSOOUT..MIDUSER1
+0020 0DD5C5E3 C14BC1F5 F0C9F0F7 F2F14040 40400000 40D481A3 85998981 93A24060 .NETA.A50I0721 .. Materials -
+0040 40D79996 978599A3 A8409686 40C9C2D4 3C505040 1D603CD1 60401DE8 3CD1F340 Property of IBM.&&; .-.J- .Y.J3
+0060 F5F6F8F4 60F1F2F3 4040C35D 40C39697 A8998987 88A340C9 C2D440C3 9699974B 5684-123 (C) Copyright IBM Corp.
+0080 40F1F9F8 F06B40F1 F9F9F04B 3CD2F040 1DE83CD3 C340C193 93409989 8788A3A2 1980, 1990..K0 .Y.LC All rights
+00A0 409985A2 8599A585 844B3CD4 40401DE8 3CD4D340 E4E240C7 96A58599 95948595 reserved..M .Y.ML US Governmen
+00C0 A340E4A2 8599A2A0 D985A2A3 998983A3 858440D9 898788A3 A240603C D550401D t Users Restricted Rights -.N&; .
+00E0 E83CD5E3 40E4A2B5 6B4084A4 97938983 81A38996 95409699 408489A2 839396A2 Y.NT Use, duplication or disclos
+0100 A4998540 9985A2A3 998983A3 ure restrict
      GMT-09/03/xx 19:29:49.847887 LOC-09/03/xx 15:29:49.753679
0HEXFORMAT AID FF FID 00 EID EFE2
+0000 00E8E380 E4E2C5D9 F1404040 E3E2D6D6 E4E30300 40D4C9C4 E4E2C5D9 F1404040 ..T.USER1 TSOOUT..MIDUSER1
+0020 0DD5C5E3 C14BC1F5 F0C9F0F7 F2F14040 40400000 85843CD6 60401DE8 3CD6F340 .NETA.A50I0721 ..ed.0- .Y.03
+0040 82A840C7 E2C140C1 C4D740E2 83888584 A4938540 C39695A3 998183A3 40A689A3 by GSA ADP Schedule Contract wit
+0060 8840C9C2 D440C396 99974B3C D7F0401D 603CD940 401D603C 5A50401D 603C5A6E h IBM Corp..P0 .-.R .-.!&; .-.!>
+0080 401DE8C5 D5E3C5D9 1D60A396 40839695 A38995A4 853C5B60 401D6040 C6F17EC8 .YENTER.-to continue.$- .- F1=H
+00A0 C5D3D73C 5B6F40C6 F27EE2D7 D3C9E33C 5B7C40C6 F37EC5D5 C43C5CC9 40C6F47E ELP.$? F2=SPLIT.$0 F3=END.*I F4=
+00C0 D9C5E3E4 D9D54040 4040C6F5 7ED9C6C9 D5C43C5C E340C6F6 7ED9C3C8 C1D5C7C5 RETURN F5=RFIND.*T F6=RCHANGE
+00E0 3C5CF240 C6F77E44 D73C5C7F 40C6F87E C4D6E605 3C5D4C40 C6F97E2E E6C1D73C .*2 F7=UP.*" F8=DOWN.)< F9=SWAP.
+0100 5DD840C6 F1F07ED3 C5C6E33C )Q F10=LEFT.
      GMT-09/03/xx 19:29:49.892763 LOC-09/03/xx 15:29:49.798555
0HEXFORMAT AID FF FID 00 EID EFE2
+0000 00E8E380 E4E2C5D9 F1404040 E3E2D6D6 E4E30300 D3C1E2E3 E4E2C5D9 F1404040 ..T.USER1 TSOOUT..LASTUSER1
+0020 0DD5C5E3 C14BC1F5 F0C9F0F7 F2F14040 40400000 5DE540C6 F1F17ED9 C9C7C8E3 .NETA.A50I0721 ..)V F11=RIGHT
+0040 40404040 C6F1F27E D9C5E3D9 C9C5E5C5 40401140 4013 F12=RETRIEVE . .
      GMT-09/03/xx 19:29:49.893415 LOC-09/03/xx 15:29:49.799207
0HEXFORMAT AID FF FID 00 EID EFE2
+0000 00E8E380 E4E2C5D9 F1404040 E3E2D6C9 D5408100 D3C1E2E3 E4E2C5D9 F1404040 ..T.USER1 TSOIN a.LASTUSER1
+0020 0DD5C5E3 C14BC1F5 F0C9F0F7 F2F14040 40400000 7D4040 .NETA.A50I0721 ..'
      GMT-09/03/xx 19:29:51.997746 LOC-09/03/xx 15:29:51.903538
IPCS PRINT LOG FOR USER USER1
+
2 16:15:43 09/03/xx

```

Figure 48. Example of TGET/TPUT trace output

The following fields appear in the TGET/TPUT trace. The first 2 bytes in each row show the hex offset in storage. The data follows that.

Byte (hex)	Meaning
00–03	ASCB address
04–0B	Job name
0C–0E	C"TSO"
0F–11	C"IN" for inbound data (TGET); C"OUT" for outbound data (TPUT)
12	TGET/TPUT option flags (See the TGET/TPUT option flags entry in Table 48 on page 649 to determine what document describes these bit definitions.)
13	TGET: return code (See Table 48 on page 649 to determine what document describes TGET return codes.) TPUT: X'00' EDIT, ASID, FULLSCREEN, or CONTROL options X'01' NOEDIT option X'02' TPG macro issued X'03' NOEDIT option specified and TGP macro issued
14–17	C"INIT" for the first 228-byte section of a PIU; C"MID" for the middle sections of a PIU; C"LAST" for the last section of a PIU
18–1F	TSO user ID
20	Length of network-qualified name
21–31	Network-qualified name
32–33	Zero
34	Start of user data

Traces provided by NCP

NCP provides several kinds of traces to record the flow of network events. Each trace occurs at a different point in the network (see Figure 35 on page 308). This allows you to follow an RU through the network and determine where in the network the RU is incorrect. (The RU could be out of sequence or lost, the data in the RU could have been changed, and so on.)

The NCP traces are:

- "Generalized PIU trace"
- "Line trace" on page 354
- "Network controller line trace (3710 only)" on page 358
- "Scanner interface trace (3720, 3725, and 3745 only)" on page 359
- "Transmission group trace" on page 359

Generalized PIU trace

The generalized PIU trace collects PIU trace data collected by the NCP. The resulting trace output shows the flow of PIUs exchanged between the NCP and its attached resources. This trace is hierarchical when started for a physical unit or a line. That is, logical units associated with the physical unit are automatically traced when traffic flows to them. Likewise, when the generalized PIU trace is started for a line, physical units and logical units associated with the line are automatically traced when traffic flows to them. When the generalized PIU trace is stopped for a physical unit, the trace is reset for all logical units associated with the physical unit, regardless of how the generalized PIU trace was started.

When to use the generalized PIU trace

Use the generalized PIU trace to trace PIU data at the NCP level and to determine whether the NCP has received or sent PIU data.

Note: VTAM can start a generalized PIU trace only for the resources that it owns. A data host, which does not own any NCPs, cannot start a generalized PIU trace for an NCP or any of the NCP's subordinate resources.

Generalized PIU trace operation

Start the generalized PIU trace with the MODIFY TRACE command. For more information on the MODIFY TRACE command for the generalized PIU trace, see *z/OS Communications Server: SNA Operation*.

Make sure that the GTF with the TRACE=USR option is active before starting this trace. See "Starting the generalized trace facility (GTF)" on page 322.

For more information on printing trace output, see "Formatting and printing trace records" on page 323.

Line trace

The line trace, a joint function of VTAM and the NCP, records the status of a line each time the NCP receives data from or sends data to that line. Although the trace is controlled by VTAM, the information in the trace records is collected by the NCP. The NCP sends the data to VTAM in a PIU. A trace type indicator in the PIU indicates whether the trace is a byte line trace (type 2 scanner) or a block line trace (type 3 scanner).

Note: If the data is not from a type 1, 2, or 3 scanner, VTAM will not process the data. You must use ACF/TAP.

The line trace collects the operating parameters of a line each time a level 2 interruption occurs on that line.

- For a 3705 communication controller with a type 2 communication scanner, a level 2 interruption occurs each time a byte of data is sent or received across the line.
- For a 3705 communication controller with a type 3 communication scanner, a level 2 interruption occurs each time an NCP buffer is filled and the buffer data is sent or received across the line.
- For a 3720, 3725, or 3745 communication controller, a level 2 interruption occurs each time a message (an entire PIU) is sent or received.

You can use the line trace only for lines attached to a communication controller and operating in network control mode.

For each 3705 communication controller, as many as eight line traces can be active at a time. For each 3720, 3725, or 3745 communication controller, a combination of eight line or scanner interface traces can be active at one time.

The number of active line traces to be allowed is specified during NCP generation (default is 2). In a cross-domain network in which the communication controller is connected to more than one host processor, the number of active traces allowed is distributed among the connected host processors on a first-come, first-served basis.

Note: As the number of active line traces increases, the system becomes less efficient.

When to use the line trace

You might use this trace if you suspect a problem with a device attached to a communication controller. If data appears correctly in the line trace but the terminal or printer does not react appropriately, the device itself is probably causing the failure.

Before using a line trace, you should use buffer and I/O traces to verify that the problem is not in VTAM or an application program. You may also want to use the scanner interface trace (for 3720, 3725, and 3745 communication controllers only), which traces data after it has been processed by the NCP and before it leaves the communication controller. Therefore, it can help determine whether the problem is in the NCP or in the line.

Note: VTAM can start a line trace only for the resources that it owns. A data host, which does not own any NCPs, cannot start a line trace for an NCP or any of the NCP's subordinate resources.

You might want to use a transmission group (TG) trace instead of or in addition to this trace. If there is more than one active line in a transmission group, and you do not know which line is causing the problem, use the transmission group trace. The transmission group trace also shows more data than the line trace.

A line trace can be active for any line in the transmission group. However, a line trace and a transmission group trace cannot be active for the same line, at the same time.

Line trace operation

Start the line trace with the MODIFY TRACE command or the TRACE start option. In either case, specify TYPE=LINE. For more information on the MODIFY TRACE command, see *z/OS Communications Server: SNA Operation*. For more information on the TRACE start option, see the *z/OS Communications Server: SNA Resource Definition Reference*.

Make sure that the GTF with the TRACE=USR option is active before starting this trace. See “Starting the generalized trace facility (GTF)” on page 322.

To format and print line trace data for a 3705, 3720, 3725, or 3745, use ACF/TAP and specify INPUT=LINE.

For more information on printing trace output, see “Formatting and printing trace records” on page 323.

Line trace output (CS type 2)

Figure 49 on page 356 shows an example of line trace output for a 3705 communication controller with a type 2 communication scanner.

The trace fields are explained after the figure. In addition to the fields described here, additional operating-system-dependent fields may appear. For a description of these fields, see Table 16 on page 325.

```

LINE LINE01 LRC(000,000) INBOUND ACTIVE RNTIME=1D
LCD C PCF A TIME 16 SCF 42 PDF D4 LCD C PCF A TIME 16 SCF 42 PDF 40
LCD C PCF A TIME 16 SCF 42 PDF E5 LCD C PCF A TIME 16 SCF 42 PDF 40
LCD C PCF A TIME 16 SCF 42 PDF E3 LCD C PCF A TIME 16 SCF 42 PDF 40
LCD C PCF A TIME 16 SCF 42 PDF C1 LCD C PCF A TIME 16 SCF 42 PDF 40

```

Figure 49. Example of line trace output (CS type 2)

The following fields appear in the line trace for CS type 2.

The header portion of the line trace record contains these fields:

Field header	Field contents
LINE <i>linename</i>	The name of the node being traced. In Figure 49, the line name is LINE01.
LRC(<i>xxx,yyy</i>)	The number of records lost since the last trace record was written because the trace facility could not get a VTAM buffer. <i>xxx</i> is the destination's lost record count, and <i>yyy</i> is the source's lost record count.
INBOUND	The direction of the data with respect to this host. It always says INBOUND because VTAM always receives the trace records from the NCP.
<i>status</i>	<p>The status of the line being traced. In Figure 49, the line status is ACTIVE. The line status may also appear as DEACTIVATE or SLOWDOWN.</p> <ul style="list-style-type: none"> • ACTIVE means that the line trace is active. • DEACTIVATE means that the line trace is not active. • SLOWDOWN means that the line trace is not active because the NCP is in slowdown mode. <p>A status of DEACTIVATE or SLOWDOWN appears only in the last record to be sent for that line. It means that no more data will be sent until the trace is activated again for that line.</p>
RNTIME= <i>hh</i>	A timer field. <i>hh</i> is a hexadecimal value indicating in tenths of a second the time at which the communication controller sent the completed line trace record to VTAM. This value is taken from a timer that is reset to 0 after 25.5 seconds. This value can be compared with the TIME value for the first level 2 interruption data contained in the second part of the line trace record. This comparison shows the elapsed time covered by this trace record. Comparing this with the number of level 2 interruptions contained in the trace record indicates the I/O activity during the elapsed time.

The data portion of the line trace record contains a timer value and values from the interface control word (ICW) and communication controller hardware registers at each level-2 interruption. Each printed line contains up to two sets of data, one for each of two level-2 interruptions. To determine what document contains detailed descriptions of the fields in line trace records, see Table 48 on page 649.

Field header	Field contents
LCD	The line control definer (LCD) indicates the type of protocol being used over the line. The LCD for a line is generally set when the NCP is loaded and remains unchanged throughout normal operation.
PCF	The primary control field (PCF) indicates the state of the line interface at any particular time.

Field header	Field contents
TIME	This hexadecimal value indicates, in tenths of a second, the time elapsed between level-2 interrupts as they occur on the line interface as data is received. Level-2 interrupts are represented by the data (PDF). The value of this timer is reset to 0 every 25.5 seconds.
SCF	The secondary control field (SCF) is used as a status and operation modifier by the communication scanner and the control program.
PDF	The parallel data field (PDF) is used as a character buffer. For a transmission, the program places the characters to be sent in the PDF, and transmits them to the line interface. For a receive operation, the character is assembled, then transferred by hardware to the PDF.

Line trace output (CS type 3)

Figure 50 shows an example of line trace output for a 3705 communication controller with a type 3 communication scanner. This trace record shows the operating status of a line each time an NCP buffer is filled, instead of each time a character is transmitted.

The trace fields are explained after the figure. Most fields in this record are similar to the example for the type 2 scanner, but the type 3 scanner has some additional fields. To determine what document contains detailed descriptions of the additional fields in CS type 3 output, see Table 48 on page 649.

```

LINE LINE01 LRC(000,000) INBOUND ACTIVE RNTIME=1D
  STATUS SCF LCD PCF EPCF STAT1 STAT2 TIME ADDR CNTL IOBLXB DISP ICW CCBLV
      44 9 9 00 00 01 80 03 71 30 AA 11 1BB2
TEXT C0B4

```

Figure 50. Example of line trace output (CS type 3)

In addition to several fields shown in the line trace for CS type 2, the following fields appear in the line trace for CS type 3:

Field header	Field contents
EPCF	The extended primary control field. These bits extend the meaning of the PCF.
STAT1	Status byte 14 of the ICW.
STAT2	Status byte 15 of the ICW.
TIME	This hexadecimal value indicates, in tenths of a second, the time elapsed between level-2 interrupts as they occur on the line interface as data is received. Level-two interrupts are represented by the data. The value of this timer is reset to 0 every 25.5 seconds.
ADDR	The SDLC frame address field.
CNTL	The SDLC frame control field.
IOBLXB	The 1-byte command field from the IOB/LXB control block.
DISP	DISP is byte 0 from the communication scanner display register.
ICW	Byte 13 of the interface control word.
CCBLV	CCBLV is the address of the level 2 interruption processor routine, taken from field CCBL2 in the NCP CCB control block.

Field header	Field contents
TEXT	The contents of the NCP buffer being transmitted in both hexadecimal and EBCDIC.

Network controller line trace (3710 only)

The network controller line trace facility traces SDLC, BSC, and start/stop data link control frames sent or received by an IBM 3710 Network Controller.

The 3710 physical unit provides a network controller line trace on a physical unit type 2 node. The network controller line trace facility traces data link control (DLC) protocols, such as DLC frames exchanged between a network controller and an NCP, or DLC frames exchanged between a network controller and a control unit.

The network controller line trace facility also traces the synchronous data link control (SDLC), binary synchronous control (BSC), and start/stop DLC frames that are sent or received by a line adapter.

When tracing an SDLC or BSC line, you can choose to trace all the frames (control and data) or just the data frames. All frames are traced on a start/stop line.

Note: If the data is not from a type 1, 2, or 3 scanner, VTAM will not process the data. You must use ACF/TAP.

When to use the network controller line trace

Run this trace if you have an incorrect output problem, a performance problem, or if you suspect a network controller problem.

Network controller line trace operation

Start the network controller line trace with the MODIFY TRACE command. For more information on the MODIFY TRACE command for the network controller line trace, see *z/OS Communications Server: SNA Operation*.

Active traces stop if:

- The network controller is reset (for example, during a failure or a configuration load).
- The network controller enters slowdown (shown in the header of the RECTRD RU).
- The line fails during the trace.

If the network controller recovers after a failure, and it is not the only one on the line, the trace will resume.

Make sure that the GTF is active before starting this trace. See “Starting the generalized trace facility (GTF)” on page 322.

Use ACF/TAP with NETCTLR specified to format and print the output. For more information on printing trace output, see “Formatting and printing trace records” on page 323.

Network controller line trace output

VTAM receives the trace data from a network controller in a series of RECTRD RUs, which are put into trace records. Use ACF/TAP to print the trace output.

Scanner interface trace (3720, 3725, and 3745 only)

The scanner interface trace, a joint function of VTAM and the NCP, stores the operating parameters of a line whenever a 3720, 3725, or 3745 communication controller sends or receives a message. Although the trace is controlled by VTAM, the information in the trace records is collected by the NCP. The scanner interface trace collects the operating status of a line whenever the communication controller sends or receives a message (PIU).

When you start the trace, you can choose how many bytes of data you want to trace. The default is the entire PIU. You can use the scanner interface trace only on lines operating in network control mode.

For each communication controller, a combination of eight line traces or scanner interface traces can be active at a time. The number of active scanner interface traces to be allowed is specified during NCP generation (default is 2). In a multisystem network in which the communication controller is connected to more than one host processor, the number of active traces allowed is distributed among the connected host processors on a first-come, first-served basis.

Note: As the number of active scanner interface traces increases, the system becomes less efficient.

When to use the scanner interface trace

Use the scanner interface trace to determine whether the NCP or a line is causing a problem. The scanner interface trace collects inbound data before the NCP processes it, and collects outbound data after the NCP processes it. Therefore, if outbound data is correct in a scanner interface trace, but the device does not react properly, the problem is in either the line or the device.

Scanner interface trace operation

Start the scanner interface trace with the MODIFY TRACE command or the TRACE start option. The COUNT option allows you to choose how many bytes of data you want traced (0 to 254 or ALL). The default is the entire PIU.

For more information on the MODIFY command for the scanner interface trace, see *z/OS Communications Server: SNA Operation*. For more information on the TRACE start option, see *z/OS Communications Server: SNA Resource Definition Reference*.

Make sure that the GTF with the TRACE=USR option is active before starting this trace. Use ACF/TAP with INPUT=LINE to format and print the output.

For more information on printing trace output, see “Formatting and printing trace records” on page 323.

Transmission group trace

The transmission group (TG) trace, a joint function of VTAM and the NCP, traces the sequence of PIUs being sent through a transmission group. The transmission

group trace shows PIU traffic on a transmission group as though the transmission group were a single line. The sequence of PIUs traced is the sequence of their entry to and exit from the transmission group, not the sequence of actual transmission along the physical lines in the transmission group. The line trace shows the sequence of PIUs transmitted and received on a physical line.

When to use the transmission group trace

Use the transmission group trace instead of a line trace if your failure is restricted to sessions using a specific transmission group and you have more than one active line in a TG.

Note: VTAM can start a transmission group trace only for the resources that it owns. A data host, which does not own any NCPs, cannot start a transmission group trace.

Transmission group trace operation

Start the transmission group trace with the MODIFY TRACE command. The name specified in the ID operand is any line associated with the transmission group trace. The trace is started for the transmission group in which the specified line resides. This line and its associated link station must be active before the transmission group trace is started. If a line trace is already active for the chosen line within a transmission group, you must stop it before starting a transmission group trace.

When activated, the transmission group trace remains active until one of the following conditions occurs:

- The operator stops the trace.
- The associated line or link station is deactivated or fails.
- The NCP goes through automatic network shutdown.
- The NCP goes into slowdown mode.

If the transmission group trace is ended because the associated line or link station fails or is deactivated, the operator can restart it by issuing the MODIFY command for another line in the TG.

For more information on the MODIFY TRACE command for the transmission group trace, see *z/OS Communications Server: SNA Operation*.

Note: If the data is not from a type 1, 2, or 3 scanner, VTAM will not process the data. You must use ACF/TAP.

Make sure that the GTF with the TRACE=USR option is active before starting this trace. See “Starting the generalized trace facility (GTF)” on page 322.

Use IPCS or ACF/TAP to format and print these trace records. For IPCS, specify USR(LINE) or USR(FF2) on the GTFTRACE option. For ACF/TAP, specify INPUT=LINE.

For more information on printing trace output, see “Formatting and printing trace records” on page 323.

Transmission group trace output

Figure 51 on page 361 is an example of transmission group trace records.

The trace fields are explained after the figure. In addition to the fields described here, additional operating system-dependent fields may appear. For a description of these fields, see Table 16 on page 325.

```

LINE  LINE23  LRC(000,000)  INBOUND  ACTIVE  RNTIME=00
TEXT  42000000  2000009A  00000004  0000000C  0E00000B
      0005808D  000ACB80  00A0
TEXT  6E000000  F000009A  00000004  00000001  00000000
      00000032  20000000  00000823  00000000  00000000
      00000123  00000040

```

Figure 51. Example of transmission group trace output

Table 21 describes the fields in the transmission group trace.

Table 21. Fields in the transmission group trace

Field header	Meaning
LINE <i>linename</i>	The name of the line associated with the transmission group trace. This line name was specified in the MODIFY TRACE command that started the transmission group trace. The line name is LINE23 in Figure 51.
LRC(<i>xxx,yyy</i>)	The number of records lost since the last trace record was written because the trace facility could not get a VTAM buffer. The destination's lost record count is <i>xxx</i> , and the source's lost record count is <i>yyy</i> .
INBOUND	The direction of the data with respect to this host. It always says INBOUND because the data is always received from the NCP.
<i>status</i>	The status of the line being traced. In Figure 51, the line status is ACTIVE. The line status may also appear as DEACTIVATE or SLOWDOWN. <ul style="list-style-type: none"> ACTIVE means that the line trace is active. DEACTIVATE means that the line trace is not active. SLOWDOWN means that the line trace is not active because the NCP is in slowdown mode. <p>A status of DEACTIVATE or SLOWDOWN appears only in the last record to be sent for that line. It means that no more data will be sent until the trace is activated again for that line.</p>
RNTIME= <i>hh</i>	A timer field, where <i>hh</i> is a hexadecimal value indicating, in tenths of a second, the time at which the communication controller sent the completed line trace record to VTAM. This value is taken from a timer that is reset to 0 every 25.5 seconds. This value can be compared with the TIME value for the first level 2 interruption data contained in the second part of the line trace record. This comparison shows the elapsed time covered by this trace record. Comparing this with the number of level 2 interruptions contained in the trace record indicates the I/O activity during the elapsed time.
TEXT	Shows the TH and RH for each PIU that traversed the data path through the transmission group. If the PIU is for data flow control, session control, or network control, the full RU portion of the PIU is also included in the trace record. The RU may contain sense data. If the PIU is function management data (FMD) and contains an FM header, six bytes of the RU are included in the transmission group trace record. If an FMD PIU without an FM header is traced, the RU is not included.

Traces provided by TCP/IP

The TCP/IP program provides multiple kinds of traces to trace TCP/IP problems. For SNA Enterprise Extender (EE) traffic, the packet trace and data trace facilities in TCP/IP may prove beneficial, as they allow both the collection and formatting of this type of SNA traffic.

The OSAENTA command facility in TCP/IP might be beneficial for debugging other network problems, including VTAM and SNA problems. The VARY TCPIP,OSAENTA command provides the ability to trace data flowing over the PCI bus in an OSA-Express2 or later adapter configured in QDIO mode, whether the data is flowing to or from the network, TCP/IP, Enterprise Extender, or Linux. This trace facility enables you to determine whether the data flowing outbound to the adapter reached the network, or whether the data flowing inbound from the network reached the adapter. See z/OS Communications Server: IP Diagnosis Guide for details about how to enable and format the OSAENTA trace.

```

1010 SLOVAKIA OSAENTA 00000007 20:43:11.637547 OSA-Express NTA
To Interface      : EZANTA0GETHF                               Full=240
Tod Clock        : 2006/10/18 20:43:11.637547
Frame: Device ID : N/A                               Sequence Nr: 1462       Discard: 0 (OK)
Segment #       : 0                                       Flags: Tunnel Out Nta L3
Source          : ::0
Destination     : ::0
Source Port     : 0                                       Dest Port: 0           Asid: 0000 TCB: 00000000
IpHeader: Version : 6                                       Header Length: 40
Trafcls        : 00                                       Dscp: 00 (CS0)        ECN: 00 (Default)
Payload Length  : 182                                       Flow: 000000
Hops           : 64                                       Protocol: UDP
Source         : 2000:197:11:116::1
Destination    : 2000:197:11:115::1

UDP
Source Port     : 12001 (EE-Network) Destination Port: 12001 (EE-Network)
Datagram Length : 182                                       CheckSum: 15E1 FFFF

Ip Header       : 40                                       IP: 2000:197:11:116::1, 2000:197:11:115::1 Offset: 0
000000 60000000 00B61140 20000197 00110116 00000000 00000001 20000197 00110115
000020 00000000 00000001

Protocol Header : 8                                       Port: 12001, 12001    Offset: 28
000000 2EE12EE1 00B615E1

Data           : 174   Data Length: 174                               Offset: 30
000000 280403C6 08D40000 00000000 00FF0012 C0C58100 0105433C 04000800 00007F00 |...F.M.....{Ea.....".|
000020 00063603 22853000 1B382000 0000005C 00000900 00000000 0000020B 91812905 |.....e.....*.....ja..|
000040 02FF0003 D0000004 22F0F0F3 001910D5 C5E3E2D6 E4E3C84B D9D6D4C1 D5C9C100 |....}...003...NETSOUTH.ROMANIA.|
000060 00000000 00000000 4712C440 00001013 1C60D723 E3B459B0 6CF011D5 C5E3E2D6 |.....D .....-P.T...%0.NETSO|
000080 E4E3C84B E2D3D6E5 C1D2C9C1 06810000 01131C35 08900040 900212C4 10D5C5E3 |UTH.SLOVAKIA.a..... ...D.NET|
0000A0 E2D6E4E3 C84BD9D6 D4C1D5C9 C100                                SOUTH.ROMANIA.

Padding, FCS    : 18
000000 20000197 00110123 00000000 00000003 7C00                                |...p.....@.|

```

Figure 52. Sample TCP/IP trace of EE data (part 1 of 2)

```

Encapsulation      : 1                Offset: 30
LLC: Dsap(I)      : 28 ( )            Ssap(C): 04 (SNA)
Unnumbered(P)    : 03 (UI)
NLH Anr Route
TpF               : Network          Flags: No_Delay
Type              : ANR Label        TP          ER Number  Address
NCE               : D4000000 00000000 N/A          N/A          N/A
Thdr
TCID              : 12C0C581 00010543
Reuse_Ct         : 12C0C581          Index: 0001          Element: 0543
Flag1            : 3C04 (SMSG EMSG STRQ REPLY OPTS)
Offset           : 32                Length: 127          Sequence: 1590
Segment 22: Size : 12                Offset: 14           Adaptive RB Pacing
Flags            : 85 (Req Normal)    Rate: 3              Reply: 0
Field1           : 1783840            Field2: 0
TH Version       : 5                  Flags: RHI CMPLI     Sequence: 9
Session Addr     : 00000000 00000002
Rh - FMH Request : FMH-5 (Attach)
RH               : 0B9181 (FI BCI ECI DR1 ERI PI BBI CEBI)
FMH5: Lenth      : 41                Type: 02FF           Flag2: 00 ( )
RscTp           : BASIC               Flag3: 00 ( )
TpNam(4)        : 22F0F0F3           'FMH-5 Attach-Locate'
FQNAM(16)       : NETSOUTH.ROMANIA    LUOW: 000000000000.0000

GDS 12C4: Len    : 71                Offset: 29           GDS Locate
Flags           : 40                Srch_Num: 4115
Control Vectors:
CV60: Len       : 28                Offset: 32           Fully Qualified PCID
PCID            : D723E3B459B06CF0 CP: NETSOUTH.SLOVAKIA
CV81: Len       : 6                 Offset: 4E           Context specific
00 06810000 0113                *.a....             *
CV35: Len       : 28                Offset: 54           Extended Sense Data
Sense           : 08900040          Flags: 90
RESOURCE NOT FOUND ON BROADCAST
RUID            : 12C4
PONAME          : NETSOUTH.ROMANIA
3 control vectors formatted

LLC Header      : 3                Offset: 30
000000 280403

ANR Header      : 12                Offset: 33
000000 C608D400 00000000 0000FF00

Transport Header : 32                12C0C581 00010543    Offset: 3F
000000 12C0C581 00010543 3C040008 0000007F 00000636 03228530 001B3820 00000000 |. {Ea.....".....e.....}

TH5 Header      : 12                00000000000000002    Offset: 5F
000000 5C000009 00000000 00000002

Request Header  : 3                Offset: 6B
000000 0B9181

Data            : 112          Data Length: 112          Offset: 6E
000000 290502FF 0003D000 000422F0 F0F30019 10D5C5E3 E2D6E4E3 C84BD9D6 D4C1D5C9 |.....}....003...NETSOUTH.ROMANI
000020 C1000000 00000000 00004712 C4400000 10131C60 D723E3B4 59B06CF0 11D5C5E3 |A.....D .....-P.T...%0.NET
000040 E2D6E4E3 C84BE2D3 D6E5C1D2 C9C10681 00000113 1C350890 00409002 12C410D5 |SOUTH.SLOVAKIA.a..... ..D.N
000060 C5E3E2D6 E4E3C84B D9D6D4C1 D5C9C100 |ETSOUTH.ROMANIA.

Padding, FCS    : 18
000000 20000197 00110123 00000000 00000003 7C00          |...p.....@.

```

Figure 53. Sample TCP/IP trace of EE data (part 2 of 2)

Chapter 8. Using the VIT analysis tool

This topic includes the following subtopics:

- “Setting up and running the VIT analysis tool”
- “Analyzing storage” on page 369
- “Counting request/response units (RUs)” on page 377
- “Extracting information from the VIT” on page 385
- “Using the timing options” on page 394
- “Using the I/O options” on page 397
- “Creating your own parameter data set” on page 399

You can use the VIT analysis tool to obtain information about a VTAM internal trace (VIT) that you have recorded on or transferred to an external device. The tool provides the following functions:

- Storage analysis
- Request and response unit (RU) counting
- VIT extraction

You can choose to process only the VIT records that fall within a given time range in the trace record. In addition, you can choose to:

- Add a title and a short description to the first page of each report.
- Format the output.
- Create a mini report at a specified interval.

For information on required target data sets for the tool, see *z/OS Communications Server: New Function Summary*. If you want a customized interface to be active to select the trace analysis commands of the VTAM program, see *z/OS Communications Server: New Function Summary* for information.

If you experience problems that you suspect to be related to the VIT analysis tool, see “VTAM internal trace (VIT) analysis tool problems” on page 44 for help.

Setting up and running the VIT analysis tool

These steps provide the minimum information that you need to set up and run the VIT analysis tool.

Procedure

Complete these steps to set up and run the tool:

1. Record a VIT on an external device or transfer a previously recorded VIT to an external device.
2. Set up to run the tool.
3. Create the parameters for the job.
4. Run the job.
5. Check the output.

Results

The following topics describe each step:

Step 1. Record a VIT

You must have a VIT on an external device, such as a disk or a tape, before you can use the tool.

A VIT that has been internally recorded can be copied to an external device using the VTAMMAP VITAL function. For instructions on the VTAMMAP VITAL function, see “VITAL” on page 286.

Step 2. Set up to run the tool

Create a data set

Create a data set specifying the input and output data sets and the tool program name. You may use JCL, a CLIST, or an REXX exec to create your data set.

Batch mode:

Use the sample JCL shown in Figure 54. Lowercase indicates required variable information. The actual JCL is determined by your installation. For example, A has been defined as a printer in the sample installation, and SYSOUT=A directs output to it. Similarly, the sample JCL assumes that all input data sets have been cataloged.

Modify the JCL by including appropriate DD names. Even though all DD names shown are not required for all runs, you might want to list them to avoid changing your JCL when you change parameters. The record format for the output data sets can be variable or variable blocked (RECFM=V or RECFM=VB).

```
//jobname JOB (account),'user name',etc.
//ISTRAFT1 EXEC PGM=ISTRAFT1,REGION=0K
//STEPLIB DD DSN=SYS1.MIGLIB,DISP=SHR
//SUMMARY DD SYSOUT=A,DCB=(RECFM=V,LRECL=84)
//DETAILS DD SYSOUT=A,DCB=(RECFM=V,LRECL=84)
//LOG DD SYSOUT=A,DCB=(RECFM=V,LRECL=124)
//OUTSTAN DD SYSOUT=A,DCB=(RECFM=V,LRECL=124)
//VITEXT DD SYSOUT=A,DCB=(RECFM=V,LRECL=124)
//PARM DD DSN=userid.run1.parm,DISP=SHR
//TRACE DD DSN=userid.run1.trace,DISP=SHR
```

Figure 54. Sample JCL for VIT analysis

As shown in Figure 54, the JCL contains the following DD names:

- SUMMARY is required and specifies where the output summarizing the trace is directed.
- DETAILS is required only for storage analysis and RU counting. It specifies where details of the trace analysis are directed.
- LOG is required. It specifies where VIT entries with possible errors are directed.
- OUTSTAN is required only for storage analysis of outstanding entries. It specifies where the list of outstanding GBLK, VTAL, and REQS entries is directed.

- VITEXT is required only for VIT extraction. It specifies where the VIT entries extracted from the trace are directed.

Note: Only the FORMAT and NOFORMAT output options should be used when directing VITEXT output to a printer as shown in Figure 54 on page 366. VITEXT output using TRACEFORMAT should be directed to disk or tape.

For VITEXT output using TRACEFORMAT, the record length must be the length of the TRACE record or 284, whichever is smaller.

- PARM is required and specifies the parameters to be passed to the VIT analysis tool. Parameters can be specified in-stream (in the JCL) or in a data set; do not use the PARM parameter on the EXEC statement for this purpose because of size restrictions. The PARM data set must have fixed records (can be blocked) with:
 - LRECL=80
 - RECFM=FB
- TRACE is required and specifies the input data set containing the trace to be processed. The TRACE DCB information must match the actual data set characteristics. The record format can be V, VB, or VBA.

Interactive mode:

As an alternative to running in batch mode, you may invoke the following routine to run the VIT analysis tool interactively. If you choose this method for processing, your terminal will be unavailable until processing is completed.

```

1. /*REXX*/
2. /*****/
3. /* Run the VIT analysis tool interactively. */
4. /* */
5. /* Tailor the data set names and other ALLOC options as needed for*/
6. /* each run. */
7. /* */
8. /*****/
9.
10.'ALLOC DD(PARM) DSN(run1.parm) SHR'
11.'ALLOC DD(TRACE) DSN(run1.trace) SHR'
12.'ALLOC DD(SUMMARY) DSN(run1.summary) OLD'
13.'ALLOC DD(DETAILS) DSN(run1.details) OLD'
14.'ALLOC DD(OUTSTAN) DSN(run1.outstan) OLD'
15.'ALLOC DD(LOG) DSN(run1.log) OLD'
16.'ALLOC DD(VITEXT) DSN(run1.vitext) OLD'
17.
18."CALL 'SYS1.MIGLIB(ISTRIFT1)'"
19.
20.'FREE DD(SUMMARY,DETAILS,OUTSTAN,LOG,VITEXT,PARM,TRACE) '

```

Figure 55. Sample VIT analysis tool interactive routine

The lines in Figure 55 are:

Line	Description
1	Required for a REXX EXEC.
2–8	Comments.
10–16	Data set allocations. The data sets must be preallocated.
18	Invokes the VIT analysis tool, assuming that it has been installed in the SYS1.MIGLIB load library.

Line	Description
20	Frees the data sets allocated to the DD name statements to allow the exec to run again with different data set names.

Note:

1. The DD parameters are required. The DSN parameters are optional and can be varied.
2. Return codes are not checked in this example.
3. Your user terminal will not be available while the tool is active.
4. This example is not shipped with the VTAM code and is included for information only.

Step 3. Create the parameters for the job

To create the parameters needed to analyze your VIT, perform the following steps:

1. Use the panel interface. On the VTAM Internal Trace Analysis panel, specify the form of processing you want to use. Only one function may be used in a session. The panel interface then provides choices and help in specifying values for the parameters, and creates the PARM data set.

Pick option 1, 2, or 3 from the Figure 56 ISTT0001 and follow the processing path until you return to the panel ISTT0001. Then choose option 4 to indicate that your input has been completed.

```

ISTT0001          VTAM Internal Trace Analysis

Select a choice, then press Enter.

— 1. Storage Analysis
   2. Request/response unit counting
   3. VIT extraction
   4. Input Complete

(C) Copyright IBM Corporation 1993,2002. All rights reserved.
Command ==> _____
:
:

```

Figure 56. VTAM internal trace analysis option panel

Note: If an incorrect value is entered in a field, the cursor appears on the field where the error was made. For help about that field, press F1.

2. Use an editor to either:

- Create a parameter data set.
- Code the parameters in-stream in the JCL created in step 2.

See the parameter syntax for a particular function, and “Creating your own parameter data set” on page 399 for further details.

See "How to read a syntax diagram" for general information on how to code and read syntax diagrams.

Step 4. Run the job

Submit the data set for processing that you created in step 2. You may process it in batch mode or interactively.

Step 5. Check the output

After processing is complete, check the following data sets for your results:

- For RU counting and storage analysis reports, check the DETAILS data set.
- For the extracted VIT entries, check the VITEXT data set.
- For unmatched storage allocation entries, check the OUTSTAN data set.

If you do not get the expected output, check the SUMMARY and LOG data sets for error messages or other information on what might have caused the problem. For example, the SUMMARY data set contains the parameters used for the job, including the parameters specified and the defaults taken.

If the trace has wrapped, indicate this on the I/O Options panel or use the WRAP parameter. If the GTF trace tapes were specified in the wrong order, correct the order of the tapes in the JCL. In either case, submit the job again.

Return codes

If an error has occurred and the SUMMARY data set is available, a message will be written in the data set. The return codes are:

0	No errors found		
4	Counter overflow	-	processing continues, if possible
8	Storage unavailable	-	processing continues, if possible
10	I/O failure	-	processing continues, if possible
12	Unrecoverable error	-	processing stops

Environment

Environmental factors are:

- The VTAM formatted trace and the VIT analysis tool cannot process data created by earlier VTAM releases because of changes to output formats.
- Because of the way that GTF handles entries that continue to multiple records, VTAM can only assume that the continued records are contiguous, and matching the continuation record to the prior record cannot be guaranteed.
- If you get trace information that is out of sequence, the trace may have wrapped. If the trace wrapped when it was recorded, specify WRAP on the VIT analysis tool.
- The existing VTAM formatted trace provides trace record formatting by splitting up all VIT entries into logical pieces and adding labels to indicate what the data represents. This function will not be replaced, and you can still format the VIT using the IPCS subcommands VERBEXIT VTAMMAP or GTFTRACE .

Analyzing storage

Use storage analysis to count storage allocated and freed, match related SMS and CSM entries, and report potential storage concerns found in a VIT. The main panel for storage analysis is shown in Figure 57 on page 370. Select an option and follow the prompts.

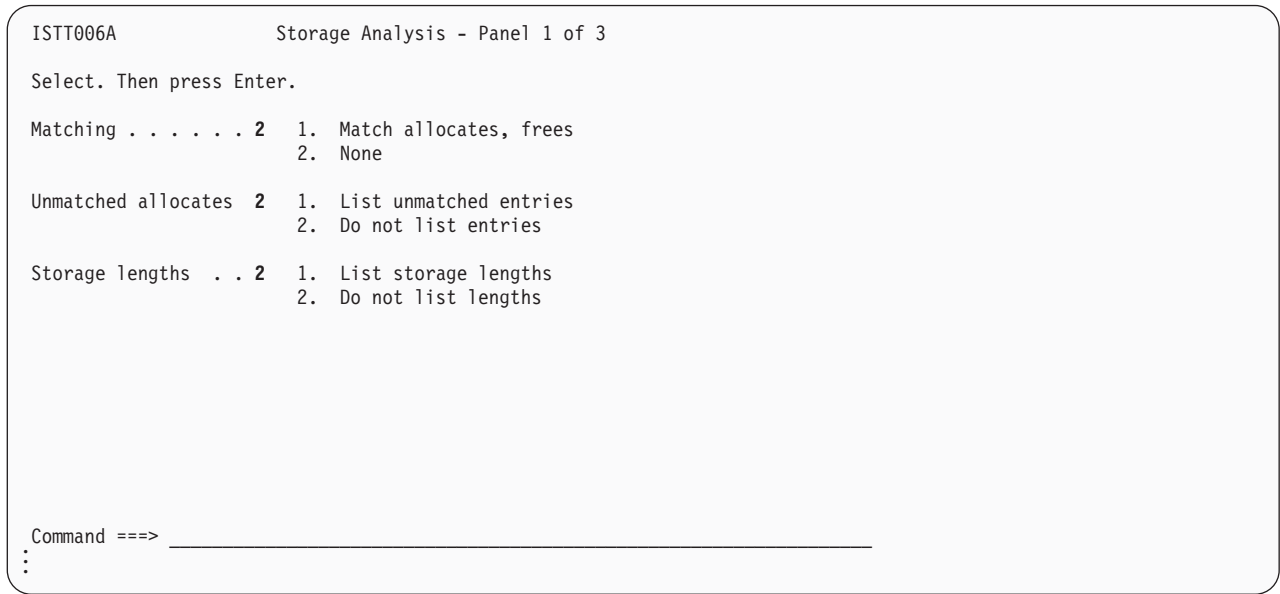


Figure 57. VTAM storage analysis option panel

Matching

You can choose to match allocate entries with free entries. Depending upon which areas of storage are selected on subsequent panels, some or all of the following VIT entries are matched:

- ASNB and FRBF
- GBLK and FBLK
- GETS and FRES
- GTBF and FRBF
- REQS and QREQ
- REQS and RELS
- REQS and AREL
- VTAL and VTFR

The default is no matching.

Unmatched allocates

You can choose to list all unmatched (outstanding) storage allocates found. By default, they are ignored. There may be a large number of unmatched storage allocates and the output may be long. Unmatched allocates can occur if the VIT is not complete. Entries listed are not necessarily error conditions.

Storage lengths

You can choose to list the storage lengths (number of bytes or buffers) requested, allocated, and freed. By default, storage lengths are not listed. This option applies only to SMS entries.

GBLK pools, GETS pools, VTAL pools, REQS buffer pools, and CSM buffers

You can designate which SMS and CSM VIT entries to process. By default, all GBLK pools, GETS pools, VTAL pools, REQS buffer pools, and CSM buffers are processed. Choices are available to process some or none of these pools. If you choose to process some, a panel is displayed from which you can choose the specific pools to process. For GBLK, GETS, and VTAL, your choices include listing only the storage pools allocated from private storage or only the storage pools allocated from CSA (by default, both private and CSA storage pools are listed).

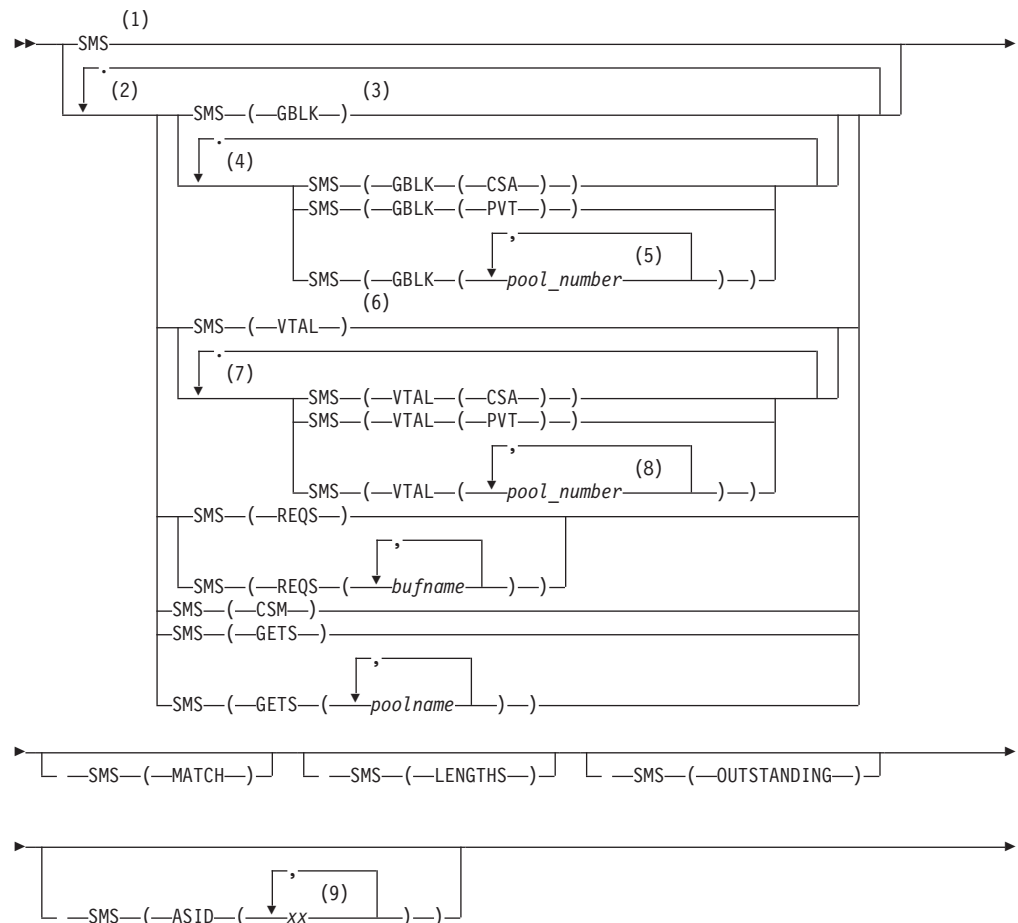
Address space identifiers (ASIDs) and data spaces


You can limit storage analysis to specific ASIDs and data spaces. By default, all ASIDs and data spaces are processed. If one or more ASIDs or data spaces are specified, only SMS entries associated with each of the specified ASIDs or data spaces are processed.

Note: Data spaces apply only to GBLK entries.

Parameter syntax

Use the following syntax only if you are using an editor to write your parameters as an alternative to the panel interface. See "How to read a syntax diagram".





Notes:

- 1 Coding SMS has the same effect as coding SMS(GBLK) SMS(VTAL) SMS(REQS) SMS(GETS) SMS(CSM).
- 2 You can code one or more from each choice group.
- 3 Coding SMS(GBLK) has the same effect as coding SMS(GBLK(PVT)) SMS(GBLK(CSA)).
- 4 You can code one or more from each choice group.
- 5 Code one or more *pool_numbers* in the range 0 - 255. Separate each *pool_number* from the next with a comma.
- 6 Coding SMS(VTAL) has the same effect as coding SMS(VTAL(PVT)) SMS(VTAL(CSA))
- 7 You can code one or more from each choice group.
- 8 Code one or more *pool_numbers* in the range 0 - 255. Separate each *pool_number* from the next with a comma.
- 9 Replace the two *xs* with 2 hex characters. You can code up to five ASIDs (machine IDs and task IDs).
- 10 Replace the eight *cs* with eight alphanumeric characters. You can code up to five DSNAMES.
- 11 Data is not case-sensitive.

Sample output for storage analysis

In this example, the following PARM file was submitted:

```
sms sms(lengths) sms(match) sms(outstanding)
stoptime(02:04:10) noformat
DESC This example shows storage analysis with a stop time.
```

The resulting Summary, Detailed, and Outstanding data sets follow.

Note:

1. Ellipses indicate that part of the output has been omitted.
2. STOPTIME is described in "Using the timing options" on page 394.
3. NOFORMAT and DESC are described in "Using the I/O options" on page 397.
4. General parameter coding information is described in "Creating your own parameter data set" on page 399.

The following example shows a Summary report for storage analysis, which is written to the data set name that you specify to receive the report. The DD name for the data set must be SUMMARY. The first line of all reports shows the VTAM level, the 20-character title (the default is Trace Analysis), and the date that the report was processed. The description, if specified, follows the title.

The next several lines of the report are the specified parameters and the defaults. These are followed by observations that are included to highlight important properties of the trace that were encountered during processing. The observations

will vary depending on the trace content and the options specified. The option choices are storage analysis, RU counting, and VIT extraction.

The trace statistics section contains the first and last time stamps, the record type count, the VIT entry count, and the VIT option count. For example, there may be 6 PIU VIT entries and 18 PIU2 entries, for a total of 24 PIU VIT option occurrences. Some VIT entries are not associated with a VIT option. If entries of this type are found, the total is listed beside **N/A** in the VIT option counts. If SNAP entries are found, the total is listed beside **?** in the VIT option counts.

VTAM V4 Trace Analysis Summary 92.325 11/20 11:28:47 LOC

This example shows storage analysis with a stop time.

Wrapped: No
Formatting: No
Interval: None
Start time: Beginning of trace
Stop time: 91.199 07/18 02:04:10.000 LOC (A43572C3 51A80000)
SMS: Yes
 ASIDs: All
 Options: MATCH LENGTHS OUTSTANDING
 GBLK: All
 DS Pools: All
 DSNAMEs: All
 GETS: All
 VTAL: All
 REQS: All
 CSM: All
RU: No
VITEXT: No

Observations

Only 27 GBLK entries were found in this VIT. The findings should be verified on a larger trace.

Only 7 VTAL entries were found in this VIT. The findings should be verified on a larger trace.

The high water mark is 392 bytes allocated by VTAL. This occurred at 02:04:07.715 LOC (record 115).

Only 18 REQS entries were found in this VIT. The findings should be verified on a larger trace.

Trace Statistics

First GTF Timestamp: 91.199 07/18 02:03:13.266 LOC (A435728D 36BFAE01)
First VIT Timestamp: 91.199 07/18 02:03:13.266 LOC (A435728D 36BFAE01)

Last VIT Timestamp: 91.199 07/18 02:04:07.716 LOC (A43572C1 240D7002)
Last GTF Timestamp: 91.199 07/18 02:04:10.866 LOC (A43572C4 2529DC02)

Summary of GTF Record Types

 1 Timestamp control records
 116 VIT records

 117 Total GTF records

Count of VIT Entry and Option Occurrences

VIT Entry Occurrences

```

33 FBLK          27 GBLK
 6 PIU           18 PIU2
19 RELS          18 REQS
 7 VTAL          5 VTFR

```

```

VIT Option Occurrences
 24 PIU
109 SMS
-----
133 Total

```

The following example shows a Detailed Report for storage analysis, which is written to the data set name that you specify to receive the report. The DD name for the data set must be DETAILS. In this example, the get block (GBLK) pool totals follow the title and description. The totals are listed first for each GBLK pool in the "Home" data space, then for each pool in other data spaces, if any exist. For each pool, the pool number is shown in decimal and hexadecimal, followed by the pool name, an indication of allocation from private storage or CSA, and the data space name.

After the counts of total entries, matches, bytes allocated and freed, and other entries, you will find a summary number for the allocate and freed entries in this pool for each storage size. This section is written only if the LENGTHS option is specified. In this example, there were 9 GBLK entries, each of which allocated 152 bytes and requested 152 bytes. The largest number of GBLK entries not matched to a free block (FBLK) at one time was five. There were 9 FBLK entries, each of which freed 152 bytes.

The totals for all GBLK pools are listed after the GBLK pool counts for each pool. Next, the VTAL pools are listed like the GBLK pools, and finally the REQS are listed.

```
VTAM V4      Trace Analysis      Detailed Report      92.194 07/12 18:51:44 LOC
```

This example shows storage analysis with a stop time.

GBLK Pool Totals:

Home Data Space:

```

GBLK Pool 0 (X'00'): RUPEPRIV (Private) Data space: Home
                    9 GBLK entries (including failures, if any)
                    9 FBLK entries (including failures, if any)
                    5 GBLK and FBLK matches
                    5 Largest number of GBLK entries at one time
                    1,368 Bytes allocated
                    1,368 Bytes allocated above the 16M line
                    1,368 Bytes requested
                    1,368 Bytes requested above the 16M line
                    1,368 Bytes freed
                    1,368 Bytes freed above the 16M line
                    608 Bytes not freed
                    760 Most unfreed bytes

```

Bytes Allocated	Bytes Requested	GBLK Entries	Maximum Requested	FBLK Entries
152	152	9	5	9

⋮

GBLK Totals:

```

27 GBLK entries (including failures, if any)
33 FBLK entries (including failures, if any)
16 GBLK and FBLK matches
13 Largest number of GBLK entries at one time

```

```

    5 Largest number of GBLK entries at one time in one pool
      was in pool 0 (X'00') in Home data space
4,192 Bytes allocated
3,696 Bytes allocated in private storage
  496 Bytes allocated in CSA
4,192 Bytes allocated above the 16M line
3,375 Bytes requested
2,888 Bytes requested in private storage
  487 Bytes requested in CSA
3,375 Bytes requested above the 16M line
3,672 Bytes freed
3,304 Bytes freed in private storage
  368 Bytes freed in CSA
3,672 Bytes freed above the 16M line
2,160 Bytes not freed
1,880 Bytes not freed in private storage
  280 Bytes not freed in CSA
2,824 Most unfreed bytes
      was at 02:04:07.715 LOC at record 116
1,536 Most unfreed bytes in one pool
      was in pool 32 (X'20') in Home data space

```

VTAL Subpool Totals:

Subpool 13 (X'0D'): Private

```

    3 VTAL entries (including failures, if any)
    2 VTFR entries (including failures, if any)
    2 VTAL and VTFR matches
    1 Largest number of VTAL entries at one time
288 Bytes allocated
288 Bytes allocated above the 16M line
192 Bytes freed
192 Bytes freed above the 16M line
 96 Bytes not freed
 96 Most unfreed bytes

```

Bytes Allocated	VTAL Entries	Most VTAL Entries	VTFR Entries
96	3	1	2

⋮

VTAL Totals:

```

    7 VTAL entries (including failures, if any)
    5 VTFR entries (including failures, if any)
    5 VTAL and VTFR matches
    2 Largest number of VTAL entries at one time
    1 Largest number of VTAL entries at one time in one pool
      was in pool 13 (X'0D')
1,344 Bytes allocated
  768 Bytes allocated in private storage
  576 Bytes allocated in CSA
1,344 Bytes allocated above the 16M line
  952 Bytes freed
  672 Bytes freed in private storage
  280 Bytes freed in CSA
  952 Bytes freed above the 16M line
  392 Bytes not freed
   96 Bytes not freed in private storage
  296 Bytes not freed in CSA
  392 Most unfreed bytes
      was at 02:04:07.715 LOC at record 115
  368 Most unfreed bytes in private
  296 Most unfreed bytes in CSA
  368 Most unfreed bytes in one pool
      was in pool 47 (X'2F')

```

REQS Totals by Buffer Pool:

REQS Buffer Pool IOBUF

5 REQS entries (including failures, if any)
7 RELS entries (including failures, if any)
4 RELS entries were matched
5 Buffers allocated
5 Buffers allocated above the 16M line
4 Buffers freed
4 Buffers freed above the 16M line
1 Buffers not freed
2 Largest number of buffers at one time

Buffers per Request	REQS Entries	Most REQS Entries
1	5	2

⋮

REQS Totals:

18 REQS entries (including failures, if any)
19 RELS entries (including failures, if any)
15 RELS entries were matched
18 Buffers allocated
18 Buffers allocated above the 16M line
15 Buffers freed
15 Buffers freed above the 16M line
3 Buffers not freed
5 Largest number of buffers at one time
was at 02:03:44.609 LOC at record 71
3 Largest number of buffers at one time in one pool
was in the LPBUF pool

CSM Totals:

8 GTBF entries (including failures, if any)
1 GTBF entries with error return code
5 ASNB output entries (including failures, if any)
1 ASNB output entries with error return code
5 FRBF entries (including failures, if any)
1 FRBF entries with error return code
3 GTBF entries were matched by FRBF
2 ASNB output entries were matched by FRBF
2 GTBF or ASNB buffer tokens duplicated
13 Buffers actually allocated by GTBF
13 Buffers requested for allocation by GTBF
5 Buffers actually assigned by ASNB
5 Buffers requested for assignment by ASNB
2 Buffers actually released by FRBF
3 Buffers actually freed by FRBF
6 Buffers requested to be freed by FRBF
9 Buffers not freed
12 Largest number of buffers at one time

The following example shows an outstanding report for storage analysis, which is written to the data set name that you specify to receive the report. The DD name for the data set must be OUTSTAN. After the title and description, the unmatched GBLK entries are listed by pool number for each data space. Next, the unmatched VTAL entries are listed by subpool number. Then the unmatched REQS and REQ2 entries are listed by buffer name.

Note:

1. The REQ2 entries consist of the VIT entry name (REQ2) followed by one to seven addresses. Instead of showing the actual REQ2 entry, which does not contain the data shown on the REQS entry, each buffer address is shown as if it had appeared in a REQS entry.

- The queued REQS entries that have not been matched by a QREQ entry are also listed in the OUTSTAN data set. These REQS have a buffer address of 0.

VTAM V4 Trace Analysis Outstanding Report 92.194 07/12 18:51:44 LOC

This example shows storage analysis with a stop time.

```

*****
* List of outstanding GBLK entries *
*****

Home Data Space:

GBLK Pool 0 (X'00'): RUPEPRIV (Private) Data space: Home

Outstanding GBLK at 02:03:52.888 LOC (record 89)
C7C2D3D2 0C000000 06638480 06357218 823D6340 00000098 864EBCF8 00000098 *GBLK.....d.....b.. ...qf+.8...q*

Outstanding GBLK at 02:04:01.276 LOC (record 102)
C7C2D3D2 0C000000 066388E0 06357218 823D6340 00000098 864EBCF8 00000098 *GBLK.....h\....b.. ...qf+.8...q*

Outstanding GBLK at 02:04:01.276 LOC (record 105)
C7C2D3D2 0C000000 066385C0 06357218 823D6340 00000098 8644B2B8 00000098 *GBLK.....e ....b.. ...qf.....q*

Outstanding GBLK at 02:04:03.374 LOC (record 108)
C7C2D3D2 0C000000 06638340 06357218 823D6340 00000098 864EBCF8 00000098 *GBLK.....c ....b.. ...qf+.8...q*
:
*****
* List of outstanding VTAL entries *
*****

Subpool 13 (X'0D'): Private

Outstanding VTAL at 02:04:05.615 LOC (record 112)
E5E3C1D3 0C000000 067A3FA0 0000000D 82478930 00000060 00000000 00000000 *VTAL.....: ....b.i.....-...

-----

Subpool 15 (X'0F'): Private

Every valid VTAL in this pool was matched by a VTFR.

-----
:
*****
* List of outstanding REQS entries *
* (Note: Each buffer from a REQ2 entry is listed as a separate REQS) *
*****

REQS Buffer Pool IOBUF

REQS waiting for QREQ or RELS at 02:03:25.703 LOC (record 35)
D9C5D8E2 0C000000 062EAE88 06321010 823B95F4 00010000 062DD648 00000000 *REQS.....h....b.n4.....0.....*

-----
:

```

Figure 58. Storage analysis with a stop time

Counting request/response units (RUs)

Use RU counting to list the number of each kind of RU found in a PIU. Because there are so many RUs, RU counting lets you specify which RUs you are interested

in. The main panel for counting request/response units is shown in Figure 59. Select an option and follow the prompts.

```
ISTT010A          Request/Response Unit Counting
Select, then press Enter.
Requests/Responses 3  1. Request units only
                     2. Response units only
                     3. Both
RUs . . . . . 1     1. All
                     2. Some - type codes
                     3. Some - list names
                     4. Some - list codes
Network addresses . - All
                   - From and/or to one address
                   - Between two addresses
                   - From one address to another
Sort order . . . . 1  1. Name
                     2. Frequency
:
```

Figure 59. VTAM request/response unit counting

Requests/responses

You can choose to process only request units, only response units, or both. Both are processed by default.

RUs

You can choose to process only specific RUs. Subsequent panels allow you to specify particular RUs by typing the codes, by picking the RUs from a list of RUs by name, or by picking the RUs from a list of RUs by code. By default, all RUs are processed.

Network addresses

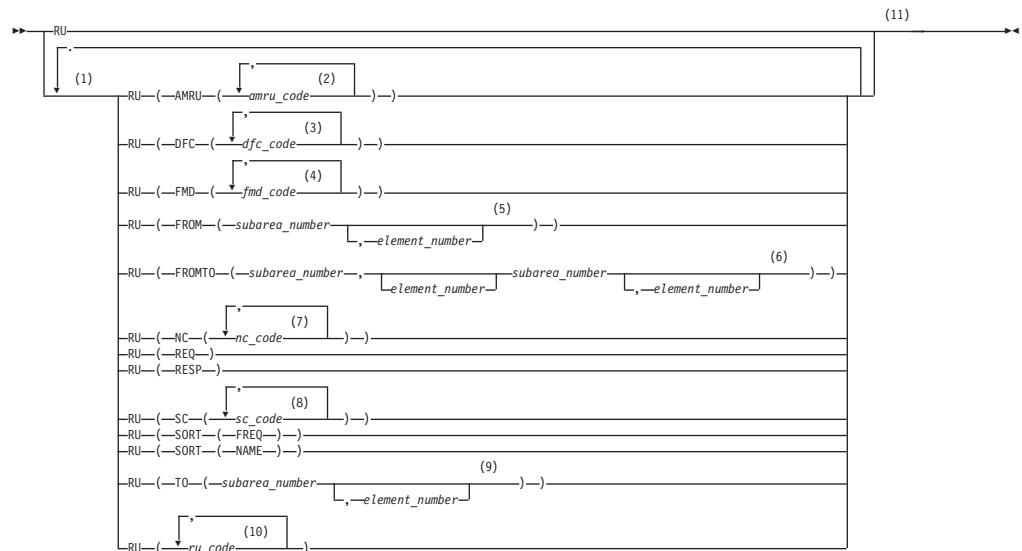
You can also specify processing of all PIUs found in the VIT regardless of the origin and destination, only PIUs from one network address to another in a single direction, all PIUs between two network addresses in both directions, or all PIUs to and from one network address.

Sorting order

You can sort the RU counts either alphabetically (by RU name) or by frequency (in descending order of counts). The default is sorting by RU name.

Parameter syntax

Use the following syntax only if you are using an editor to write your parameters as an alternative to the panel interface. See "How to read a syntax diagram".



Notes:

- 1 You can code one or more of these options. Unless otherwise noted, code each option no more than once.
- 2 Each *amru_code* is 4 or 8 hex characters. You can code up to 32 *amru_codes*.
- 3 Each *dfc_code* is 2 or 4 hex characters. You can code up to 32 *dfc_codes*.
- 4 Each *fmd_code* is 6 or 8 hex characters. You can code up to 32 *fmd_codes*.
- 5 Each *subarea_number* is 8 hex characters, and each *element number* is 4 hex characters.
- 6 Each *subarea_number* is 8 hex characters, and each *element number* is 4 hex characters.
- 7 Each *nc_code* is 2 or 4 hex characters. You can code up to 32 *nc_codes*.
- 8 Each *sc_code* is 2 or 4 hex characters. You can code up to 32 *sc_codes*.
- 9 Each *subarea_number* is 8 hex characters, and each *element number* is 4 hex characters.
- 10 Each *ru_code* is 2, 4, 6, or 8 hex characters. You can code up to 32 *ru_codes*.
- 11 Data is not case-sensitive.

RU parameter coding

RU means count all request/response unit codes.

RU(NC(...)), RU(SC(...)), RU(DFC(...)), RU(FMD(...)), or RU(AMRU(...)) means count only the RUs in categories which have the specified codes as follows:

- NC, SC, and DFC codes are 1 or 2 bytes ⁴:
 - SNA codes are 1 byte.
 - AMRU codes are 2 bytes. The first byte is X'FF'.
- FMD codes are 3 or 4 bytes ⁴:
 - SNA codes are 3 bytes.
 - AMRU codes are 4 bytes. The third byte is X'FF'.

You must fully specify all codes (the VIT analysis tool does not pad codes).

RU(...) means process only the RUs with those codes. Each code is 1, 2, 3, or 4 bytes. The VIT Analysis Tool counts these codes in any and all categories. For example, X'0D' is the NC code for NC-ACTVR and the SC code for ACTLU. If you specify **RU(0D)**, the VIT analysis tool will count both of these RU codes.

You can specify up to:

- 32 codes for **RU(NC(...))**,
- 32 codes for **RU(SC(...))**,
- 32 codes for **RU(DFC(...))**,
- 32 codes for **RU(FMD(...))**,
- 32 codes for **RU(AMRU(...))**, and
- 32 codes for **RU(...)**.

You can specify the RU codes individually or in lists. For example,

```
RU(SC)31,32,FF31,FF32))
```

is equivalent to

```
RU(SC(31)) RU(SC(32)) RU(SC(FF31)) RU(SC(FF32))
```

RU(REQ) means process requests. **RU(RESR)** means process responses. By default (if you specify neither **RU(REQ)** nor **RU(RESR)**), the VIT analysis tool counts both requests and responses.

RU(FROM(...)) means count RUs flowing from the specified network address to any network address. **RU(To(...))** means count RUs flowing to the specified network address from any network address. **RU(FROMTO(...))** means count RUs flowing from the first network address to the second network address. You can specify up to:

- 24 **RU(FROM(...))** options
- 24 **RU(To(...))** options
- 24 **RU(FROMTO(...))** options

Each subarea field (*subarea_number*) is exactly 4 bytes ⁴. Each element field (*element_number*) is exactly 2 bytes. The element fields are optional. For the **RU(FROMTO(...))** option, if you omit the first element address, keep its comma as a place-holder.

Note: You must specify the **RU(FROM(...))**, **RU(To(...))**, and **RU(FROMTO(...))** options individually. (You cannot combine them in lists like RU codes.)

By default [if you specify neither **RU(FROM(...))**, **RU(To(...))**, nor **RU(FROMTO(...))**], the VIT analysis tool ignores the origin and destination fields in the PIUs.

RU(SORT(NAME)) means sort the request and response counts by RU name (the default). **RU(SORT(FREQ))** means sort the counts by frequency (highest count first).

4. Two hexadecimal digits represent 1 byte. For example, X'FF' is 1 byte. For RU counting, you must specify all hexadecimal digits and you must omit the X and quotation marks. For example, RU(NC(X'C')) is not valid; RU(NC(0C)) is valid.

Combinations of RU options

You can combine any or all of the RU options. When considering combinations of options, you might find it helpful to think of the RU options as two groups, as shown in Figure 60 and Figure 61.

```
RU
RU(NC(...))
RU(SC(...))
RU(DFC(...))
RU(FMD(...))
RU(AMRU(...))
RU(...)
```

Figure 60. RU code options. These specify which RU codes to count.

```
RU(REQ)
RU(RES)
RU(FROM(...))
RU(TO(...))
RU(FROMTO(...))
RU(SORT(NAME))
RU(SORT(FREQ))
```

Figure 61. RU modify options. These options indicate whether to count request or response units, which origins and destinations to count, and the sort order for the count.

The options in Figure 61 modify the options in Figure 60. For example, given the following combination of options, the VIT analysis tool will count all requests for all RU codes.

```
RU(REQ) RU
```

If you specify any modify option and no code option, the VIT analysis tool uses the default, which is the **RU** option. In other words, the VIT analysis tool counts all RU codes which match the modify options. The VIT Analysis Tool prints a message in the SUMMARY data set so you will know the **RU** option is in effect.

The **RU** option (which counts all RUs) overrides the other code options (Figure 60). The VIT Analysis Tool prints a message in the SUMMARY data set so you will know the override is in effect. For example, given the following combination of options, the VIT analysis tool counts all RUs and does not check whether the specified RU, NC(04) in the example, is found in the VIT.

```
RU RU(NC(04))
```

If any of the **RU(FROM(...))**, **RU(TO(...))**, or **RU(FROMTO(...))** options match the origin or destination of an RU, the VIT analysis tool will count the RU. For example, you can ask for the counts of the following options:

- RUs from subarea X'00000012'
- RUs to subarea X'00000012'
- RUs flowing between subarea X'0000004A' element X'000C' and subarea X'00000002'

by coding:

```
RU(FROM(00000012)) RU(TO(00000012))
RU(FROMTO(0000004A,000C,00000002)) RU(FROMTO(00000002,,0000004A,000C))
```

If you list specific RUs (that is, you specify any options in Figure 60 on page 381 except **RU**), the VIT analysis tool does not count:

- User RUs
- FMH RUs
- Unknown RUs

(However, the VIT analysis tool reports the first occurrence of an unknown RU in the LOG data set.)

Sample output for RU counting

In this example, the following PARM file was submitted:

```
Desc This example shows request/response unit counting
Desc with a start time and a 30-second interval.
RU STARTIME(02:04:30) interval(00:30)
```

The resulting DETAILS, LOG, and SUMMARY data sets follow.

Note: Ellipses indicate that part of the output has been omitted. The following example shows a summary report for RU counting, which is written to the data set name that you specify to receive the report. The DD name for the data set must be SUMMARY.

This report is similar to the summary report for storage analysis. See “Sample output for storage analysis” on page 372 for details.

```
VTAM V4      Trace Analysis      Summary      92.325 11/20 11:47:26 LOC
```

```
This example shows request/response unit counting
with a start time and a 30-second interval.
```

```
Wrapped:      No
Formatting:   Yes
Interval:     00:30
Start time:   91.199 07/18 02:04:30.000 LOC (A43572D6 64780000)
Stop time:    End of trace
SMS:         No
RU:          Yes
  Options:    Requests Responses
  Sort:       Name
  Codes:     All
  From:      All
  To:        All
  From/To:   All
VITEXT:      No
```

```
*****
Observations
```

There are 7 messages in the LOG file.

Only 250 RUs were found in this trace. The findings should be verified on a larger trace.

7 responses had sense data included. See the LOG data set.

```
*****
Trace Statistics
```

```
First GTF Timestamp: 91.199 07/18 02:03:13.266 LOC (A435728D 36BFAE01)
```

First VIT Timestamp: 91.199 07/18 02:04:32.920 LOC (A43572D9 2D8DDE02)

Last VIT Timestamp: 91.199 07/18 02:06:00.891 LOC (A435732D 12B57F02)

Last GTF Timestamp: 91.199 07/18 02:06:00.891 LOC (A435732D 12B57F02)

Summary of GTF Record Types

```
      15 Timestamp control records
     6,930 VIT records
     -----
     6,945 Total GTF records
```

Count of VIT Entry and Option Occurrences

VIT Entry Occurrences

```
1,390 FBLK          6 FBL2
1,385 GBLK          6 GBL2
  104 MSG           104 MSGS
  217 MSG2          262 PIU
  746 PIU2         1,427 RELS
1,414 REQS          6 REQ2
  415 VTAL          411 VTFR
```

VIT Option Occurrences

```
  425 MSG
 1,008 PIU
 6,460 SMS
     ----
 7,893 Total
```

The following example shows a detailed report for RU counting, which is written to the data set name that you specify to receive the report. The DD name for the data set must be DETAILS. The RU counts for each interval are written after the title and description. The total count for all intervals is written at the end of the report.

Under Requests, only the first GDS variable in each RU is counted. Under Responses, all RUs with sense data included are grouped together. Each RU with sense data included is printed in the LOG data set. For example, an UNBIND response with sense data included is counted only as a response with sense data included, not as an UNBIND response.

VTAM V4 Trace Analysis Detailed Report 92.194 07/12 18:51:46 LOC

This example shows request/response unit counting with a start time and a 30-second interval.

Interval 1

First VIT timestamp in this interval:
91.199 07/18 02:04:32.920 LOC (A43572D9 2D8DDE02) (record 390)

RU Totals for Interval 1:

Requests:

```
      1 RNAA
     -----
      1 Total requests
```

Responses:

```
      1 FNA
     -----
      1 RNAA
     -----
      2 Total responses
```

```
*****
Last VIT timestamp in interval 1
91.199 07/18 02:05:01.277 LOC (A43572F4 3877DE01) (record 429)
*****
```

```
Interval 2
:
:
RU Totals:
```

Requests:

2	ACTLINK	7	BFCINIT
4	BFCLEANUP	7	BFINIT
2	BFSESSEND	5	BFSESST
5	BFTERM	18	BIND
1	BINDF	4	CINIT
5	CLEANUP	2	CONNOUT
2	CONTACT	2	CONTACTED
1	DACTLINK	1	DISCONTACT
5	FMH-5 Attach-CP Capab	6	FMH-5 Attach-TDU
9	FNA	4	GBIND BIND
4	GDS CP Capabilities	7	GUNBIND
3	INIT-OTHER	1	INOP
1	NOTIFY (SSCP<-->LU)	3	RECMS
2	REQCONT	11	RNAA
5	SESSEND	5	SESST
23	UNBIND		

157 Total requests

Responses:

2	ACTLINK	7	BFCINIT
1	BFCLEANUP	5	BIND
4	CINIT	1	CLEANUP
2	CONNOUT	2	CONTACT
1	DACTLINK	1	DISCONTACT
10	FNA	4	INIT-OTHER
1	NOTIFY (SSCP<-->LU)	11	RNAA
7	Sense Data Included	13	User
21	UNBIND		

93 Total responses

The following example shows an RU counting log, which is written to the data set name that you specify to receive the report. The DD name for the data set must be LOG. The Log contains important details found in the VIT during processing. For RU processing, the Log contains all RUs with included sense data.

VTAM V4 Trace Analysis Log 92.194 07/12 18:51:46 LOC

This example shows request/response unit counting with a start time and a 30-second interval.

Sense data included at 02:05:56.653 LOC (record 1,990)

```
Origin:      00000004 0073
Destination: 00000001 0008
Response Header: EF9000 Session Control
Sense Data:  80050000 Path error
              No session

Rejected RU code: 32
Rejected Command: UNBIND
```

Sense data included at 02:05:56.655 LOC (record 2,000)

```
Origin:      00000004 0073
Destination: 00000001 0008
Response Header: EF9000 Session Control
Sense Data:  80050000 Path error
              No session
```

```

Rejected RU code: 32
Rejected Command: UNBIND

Sense data included at 02:06:00.372 LOC (record 6,706)
Origin:          00000001 0008
Destination:    00000001 0001
Response Header: 8F9000  Function Management Data
Sense Data:     08160000 Request reject
                                   Function already inactive

Rejected RU code: 810629
Rejected Command: CLEANUP

Sense data included at 02:06:00.399 LOC (record 6,725)
Origin:          00000004 0000
Destination:    00000001 0001
Response Header: 8F9000  Function Management Data
Sense Data:     081E0001 Request reject
                                   Session reference error

Rejected RU code: 812629
Rejected Command: BFCLEANUP

Sense data included at 02:06:00.456 LOC (record 6,774)
Origin:          00000001 0008
Destination:    00000001 0001
Response Header: 8F9000  Function Management Data
Sense Data:     08160000 Request reject
                                   Function already inactive

Rejected RU code: 810629
Rejected Command: CLEANUP
:

```

Extracting information from the VIT

Use VIT extraction to extract entries from a VIT. VIT entries extracted from a VIT can be formatted, displayed in hex with the eye-catcher, or copied in the same format as the input.

Upon selecting VIT Extraction, the VIT Extraction Boolean Expression panel is displayed as shown in Figure 62 on page 386.

Type a Boolean expression or press F4 to use the template.

Operands	Description	Operators
CCcc or E'CCcc'	Option or entry name	() Delimiters
A'xxxxxxxX':nn	Address:offset	~ Not
C'CC...':nn or X'Xx...':nn	Char or Hex String:offset	- Through
B'...xxXX'	Buffer token for CSM	& And
O'xxxx...' or D'xxxx...'	Origin or Destination	Or

Command ==> _____

:

Figure 62. VIT extraction Boolean expression panel

The first time VIT extraction is invoked, the VIT extraction Boolean expression panel is blank. After the first time, the panel is displayed with the previously entered Boolean expression.

You may specify the VIT entries you want extracted by entering a Boolean expression on this panel, or by filling in a template (one or more times). See "Using the template" for information on how to use the template. Both methods result in a Boolean expression that specifies the criteria used to select VIT entries. VIT entries that contain the data specified by the Boolean expression are extracted. Extracted VIT entries may be formatted, displayed in hex with the eye-catcher, or copied as is. See "Using the I/O options" on page 397 for information on how to code these options.

Note: You cannot use the template to specify a CSM buffer token, origin, or destination for VIT extraction. See "Creating a Boolean expression without the template" on page 389 for more information.

Using the template

To use the template, press F4. The VIT Extraction Template is displayed as shown in Figure 63 on page 387.

All fields are optional. Any explicitly specified VIT entries, and VIT entries created by a specified option, are eligible for extraction if found in the trace. Only those eligible entries that meet all other specified selection criteria are extracted.

```

ISTT0014          VIT Extraction Template

Type information in one or more fields, then press Enter. This information
will be appended to the full expression.

VIT options/entries _____ +
                    _____

Address . . . . . _____ (Hexadecimal)
Offset . . . . . _____ (Decimal or Hexadecimal)

Character string . _____
Offset . . . . . _____ (Decimal or Hexadecimal)

Hexadecimal string _____
Offset . . . . . _____ (Decimal or Hexadecimal)

Command ==> _____
:

```

Figure 63. VIT extraction template

The fields on the template are described below. Fill in the template and press Enter.

The template is checked for proper data type and length of data and saved. The resulting Boolean expression is then added to the VIT Extraction Boolean Expression panel. You can append multiple instances of the template, and a VIT entry that matches any of the templates will be extracted (when VIT extraction is invoked). The length of the resulting expression is limited to the input area on the VIT Extraction Boolean Expression panel.

Press F3 to exit, and you are given the option to save the expression you have created.

The fields on the extraction template are described as follows:

VIT options/entries

You may specify VIT options or entries to limit extraction to particular VIT entries. If you specify a VIT option, VIT entries created when the designated options are active are eligible for extraction if found in the trace. For example, the LOCK option generates the LKEX, LKSH, ULKA, and UNLK trace entries. If you specify LOCK and the LOCK option was used when the VIT was started, any LKEX, LKSH, ULKA, and UNLK entries found in the trace are eligible for extraction.

Note: User-defined (SNAP) entries are allowed. For further information, see z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT.

You may also specify particular VIT entries (for example, LKEX or LKSH), and an asterisk (*) may be used in a VIT entry name to match any character (for example, LK* matches VIT entries LKEX and LKSH). If an entry with multiple parts is specified, all parts are extracted (for example, if AI2 is specified, AI1, AI2, and AI3 are extracted).

If no VIT options or entries are specified, all VIT entries are eligible for extraction.

Address

Entries with this address are eligible for extraction. Specify up to 8 hex characters for an address. If fewer than eight digits are specified, the address is padded on the left with zeros. The low-order 31 bits of the address are then used to compare on all word boundaries if an offset is not specified. If an offset is specified, then only the offset is checked for a match. Address X'00000000' is allowed only when an offset for the address is also provided.

Address offset

The offset for an address is a word boundary offset into a trace record where a comparison should be made for the address. The offset must be one of the following values:

X'04', 4
X'08', 8
X'0C', 12
X'10', 16
X'14', 20
X'18', 24
X'1C', 28

Character string

Entries containing this character string are eligible for extraction. Enter a search string of 1–30 alphanumeric characters, which includes special characters (except a single quotation mark). Data entered is not converted to uppercase. Data is case-sensitive. By default, a comparison is made at all offsets.

Character string offset

You may include a byte offset into the VIT entries where comparisons should be made. The offset can be decimal or hex. The offset range is 2–31 or X'02'–X'1F'.

Note: The length of the character string determines the offset allowed. For example, if the character string entered consists of 8 characters, the valid offset range is 2–24. The string cannot start after byte 24, because a VIT entry is 32 bytes long (byte 0 through byte 31) and 8 bytes are needed to represent the string (bytes 24 – 31).

Hexadecimal string

Entries containing this hex string are eligible for extraction. Enter a search string of 2–60 hex characters representing 1–30 bytes of data. By default, a comparison is made at all offsets.

If you code an odd number of hex digits, they are padded to the left with a 0 to make 1 byte. For example; X'D' is equivalent to X'0D'.

Hexadecimal string offset

You may include an offset into the VIT entries where comparisons should be made. The offset can be decimal or hex. The offset range is 2–31 or X'02'–X'1F'.

Note:

1. The length of the hex string divided by 2 determines the offset allowed, because two hex digits represent one byte. For example, if the hex string entered consists of 8 hex digits, 4 bytes are needed to store the string, and the valid offset range is 2 – 28. The string cannot start after

byte 28, because a VIT entry is 32 bytes long (byte 0 through byte 31) and 4 bytes are needed to represent the string (bytes 28 – 31).

2. Character and hex strings will not be found if they cross VIT entry boundaries. If a PIU is represented in the VIT as a PIU entry plus a series of PIU2 entries, and a string is requested that spans the end of one PIU2 and the beginning of a second PIU2, it will not be found.

Creating a Boolean expression without the template

The template allows you to create many Boolean expressions, but there is no way to do the following actions:

- Negate an operand
- Group operands
- Specify AND or OR
- Specify an operand multiple times on one template
- Specify a range

To do any of this, you will need to enter an expression directly on the VIT Extraction Boolean Expression panel. Blanks are recommended between all operands and operators. Again, the length of the expression is limited to the input area provided. No syntax error checking is performed.

A sample free-form expression follows:

```
AI* | (LOCK | E'PIU' & X'31':15 |  
GBL* & A'476C' - A'4800')
```

This example selects each VIT entry that is either AI1, AI2, or AI3, or fulfills at least one of the following criteria:

- LOCK option group (LKEX, LKSH, ULKA, UNLK)
- PIU entry with value hex 31 at offset 15
- GBLK or GBL2 entry with any address from X'0000476C' through X'00004800'

Operands

The following operands are used in creating a Boolean expression. For further information on these operands, see “Using the template” on page 386.

VIT options or entries

VIT entries with names identical to option names must be prefaced with an E and enclosed in single quotation marks (for example, E'MSG'). This distinguishes the name as an entry rather than an option. A VIT option or unique entry name requires no preface.

Address

Preface an address with an A and enclose it in single quotation marks (for example, A'xxxxxxx', where xxxxxxxx is the hex address).

Address offset

Specify the offset for an address after the address string, and separate the address string and the offset with a colon (for example, A'xxxxxxx':nn, where xxxxxxxx is the hex address and nn is the offset).

Note: Only decimal offsets are allowed.

Character string

Preface a character string with a C and enclose it in single quotation marks (for example, C"cccc...", where cccc... represents the character string). Data is case-sensitive.

Character string offset

Specify the offset for a character string after the character string, and separate the character string and the offset with a colon (for example, C"cccc...":nn, where cccc... is the character string and nn is the offset).

Note: Only decimal offsets are allowed.

Hexadecimal string

Preface a hex string with an X and enclose it in single quotation marks (for example, X'xxxxx...', where xxxxx... represents the hex string).

Hex string offset

Specify the offset for a hex string after the hex string, and separate the hex string and the offset with a colon (for example, X'xxxxx...':nn, where xxxxx... is the hex string and nn is the offset).

Note: Only decimal offsets are allowed.

CSM buffer token

Enter 1–12 bytes of hexadecimal digits representing a CSM buffer token. If fewer than 12 bytes are supplied, the buffer token is padded on the left with zeros. In a CSM VIT record, there may be one or two buffer tokens. The following CSM VIT records have one or two buffer tokens:

VIT Record**Number of tokens**

ASN2 2

CHG2 2

CPY3 1

CPY4 1

FIX2 2

FRB2 2

GTB3 1

PAG2 2

XBA2 1

XBI2 1

If the buffer token matches a token in a VIT entry, the VIT entry and its related entries are extracted. For example, if a GTB3 entry is matched, the corresponding GTBF, GTB2, and other GTB3 entries are also extracted.

Note: When extracting VIT entries using the CSM buffer token, bit 0 in the token is masked. Therefore, the token fields in the extracted VIT entries may not exactly match the input token.

Origin

To extract PIU entries originating at a designated network address, enter 1–12 hex characters (representing the 6-byte network address) in the form O'xx...!.

Note: The address is right-aligned and padded with zeros on the left. For example, O'3001A' means subarea 3 element 1A.

Destination

To extract PIU entries destined for a particular network address, enter 1–12 hex characters (representing the 6-byte network address) in the form D'xx...':

Note: The address is right-aligned and padded with zeros on the left. For example, D'4E' means subarea 0 element 4E.

All Specify ALL to extract all VIT entries. ALL is not valid with any other operand or operator.

Note: You may also extract all VIT entries without entering the ALL operand. For example, 'gblk | ¬ gblk' will extract all VIT entries.

Operators

The operators used in creating a Boolean expression are shown in the following table.

Table 22. Boolean expression operators in order of precedence

Operator	Description
()	Parentheses
¬	Not
-	Through
&	And
	Or

Parentheses have the highest precedence and can be used to change the normal order of evaluation. The maximum nesting level is 15. The *through* operator (a hyphen) specifies a range and can be used for addresses or a hex string in the following combinations:

- address-address
- address:offset-address
- hex string-hex string
- hex string:offset-hex string

If an offset is specified and you are using the *through* operator, the offset on the first operand is used for both operands.

Parameter syntax

Use the following syntax only if you are using an editor to write your parameters as an alternative to the panel interface.

▶▶—VITEXT— —Boolean_expression—▶▶

Note: Up to 15 VITEXT parameters may be coded for longer Boolean expressions.

VITEXT must be the first six characters and must be followed by a blank. The rest of the line is assumed to be the expression.

Sample output for VIT extraction

In this example, the following PARM file was submitted:

```
desc This example shows VIT extraction.  
desc All PIU VIT entries to or from network address 000000040073  
desc and all MSGs with the string ACTIVE will be extracted.  
vitext o'40073' | d'40073' | (MSG & c'ACTIVE')  
noformat
```

The resulting VITEXT and SUMMARY data sets follow.

Note: Ellipses indicate that part of the output has been omitted. The following example shows a Summary report for VIT extraction, which is written to the data set name that you specify to receive the report. The DD name for the data set must be SUMMARY.

This report is similar to the Summary report for storage analysis. See “Sample output for storage analysis” on page 372 for details.

This example shows VIT extraction.
 All PIU VIT entries to or from network address 000000040073
 and all MSGs with the string ACTIVE will be extracted.

```

Wrapped:      No
Formatting:   No
Interval:     None
Start time:   Beginning of trace
Stop time:    End of trace
SMS:          No
RU:           No
VITEXT:       Yes
              o'40073' | d'40073' | (MSG & c'ACTIVE')
    
```

 Observations

25 GTF VIT records were extracted and written to VITEXT.

 Trace Statistics

```

First GTF Timestamp: 91.199 07/18 02:03:13.266 LOC (A435728D 36BFAE01)
First VIT Timestamp: 91.199 07/18 02:03:13.266 LOC (A435728D 36BFAE01)

Last VIT Timestamp:  91.199 07/18 02:06:00.891 LOC (A435732D 12B57F02)
Last GTF Timestamp:  91.199 07/18 02:06:00.891 LOC (A435732D 12B57F02)
    
```

Summary of GTF Record Types

```

          16 Timestamp control records
       7,318 VIT records
       -----
       7,334 Total GTF records
    
```

Count of VIT Entry and Option Occurrences

```

VIT Entry Occurrences
  1,504 FBLK           6 FBL2
  1,498 GBLK           6 GBL2
    110 MSG            110 MSGS
    230 MSG2           279 PIU
    795 PIU2           1,464 RELS
  1,451 REQS           6 REQ2
    443 VTAL           439 VTFR
    
```

```

VIT Option Occurrences
    450 MSG
   1,074 PIU
   6,817 SMS
   -----
   8,341 Total
    
```

Figure 64. Example of VIT extraction

The following example shows a VIT Selections report, which is written to the data set name that you specify to receive the report. The DD name for the data set must be VITEXT. This report contains the VIT entries selected by the Boolean expression.

The date line containing the date, time, and record number for the first record extracted is written after the title and description. The title and description are not

written if the TRACEFORMAT option is selected. The date line is written at the beginning of the report and when the date changes. The time is written to the left of each record.

In this example, each VIT entry is written on one line and each line contains the time, the entry in hexadecimal, and the EBCDIC translation of the entry.

VTAM V4 Trace Analysis VIT Selections 92.194 07/12 18:51:50 LOC

This example shows VIT extraction.
 All PIU VIT entries to or from network address 000000040073
 and all MSGs with the string ACTIVE will be extracted.

```

*** DATE *** 91.199 07/18 02:05:56.137 LOC (A4357328 8A0BB802) (record 881)
02:05:56.137 D7C9E440 0C990000 06321010 40007870 20000037 00000001 00000004 1D000008 *PIU .r.....*
02:05:56.137 D7C9E4F2 00730018 00816B80 00310013 07B0B050 B33F8797 97870706 02000000 *PIU2.....a,.....&..gppg.....*
02:05:56.137 D7C9E4F2 00000000 00230000 04C1F0F2 D51F0008 02C3D7E2 E5C3D4C7 090300E3 *PIU2.....A02N....CPSVCMG...T*
02:05:56.137 D7C9E4F2 F9560FFA BBA50A04 D5C5E3C1 4BC1F0F2 D50004C1 F0F1D50E 0AF3D5C5 *PIU29....v..NETA.A02N..A01N..3NE*
02:05:56.137 D7C9E4F2 E3C14BC1 F0F2D50E 0AF4D5C5 E3C14BC1 F0F2D52C 0A010840 40404040 *PIU2TA.A02N..4NETA.A02N....*
02:05:56.137 D7C9E4F2 40404060 12E7E3F9 560FFABB A509D5C5 E3C14BC1 F0F2D500 00000000 *PIU2 -.XT9....v..NETA.A02N.....*
02:05:56.137 D7C9E440 0C990000 06321010 40007870 20000037 00000001 00000004 1D000008 *PIU .r.....*
02:05:56.137 D7C9E4F2 00730018 00816B80 00310013 07B0B050 B33F8797 97870706 02000000 *PIU2.....a,.....&..gppg.....*
02:05:56.137 D7C9E4F2 00000000 00230000 04C1F0F2 D51F0008 02C3D7E2 E5C3D4C7 090300E3 *PIU2.....A02N....CPSVCMG...T*
02:05:56.137 D7C9E4F2 F9560FFA BBA50A04 D5C5E3C1 4BC1F0F2 D50004C1 F0F1D50E 0AF3D5C5 *PIU29....v..NETA.A02N..A01N..3NE*
02:05:56.137 D7C9E4F2 E3C14BC1 F0F2D50E 0AF4D5C5 E3C14BC1 F0F2D52C 0A010840 40404040 *PIU2TA.A02N..4NETA.A02N....*
02:05:56.137 D7C9E4F2 40404060 12E7E3F9 560FFABB A509D5C5 E3C14BC1 F0F2D500 00000000 *PIU2 -.XT9....v..NETA.A02N.....*
:
02:05:56.712 D4E2C7E2 0C000000 00000000 0000F0F1 0039E000 C9E2E3F1 F0F5C940 C1F0F4D7 *MSGs.....01..\IST105I A04P*
02:05:56.712 D4E2C7F2 E4C3C1F3 40D5D6C4 C540D5D6 E640C9D5 C1C3E3C9 E5C56B40 D5D6C4C5 *MSG2UCA3 NODE NOW INACTIVE, NODE*
02:05:56.722 D4E2C7E2 0C000000 00000000 0000F0F1 0038E000 C9E2E3F1 F0F5C940 C1F0F4D3 *MSGs.....01..\IST105I A04L*
02:05:56.722 D4E2C7F2 D5C3C1F3 40D5D6C4 C540D5D6 E640C9D5 C1C3E3C9 E5C56B40 D5D6C4C5 *MSG2NCA3 NODE NOW INACTIVE, NODE*
02:05:56.804 D4E2C7E2 0C000000 00000000 0000F0F1 002BE000 C9E2E3F0 F9F3C940 C1F0F4E2 *MSGs.....01..\IST093I A04S*
02:05:56.804 D4E2C7F2 F1F640C1 C3E3C9E5 C56B40D5 D6C4C540 E3E8D7C5 407E40D3 C9D5C508 *MSG216 ACTIVE, NODE TYPE = LINE.*
02:05:56.883 D4E2C7E2 0C000000 00000000 0000F0F1 002FE000 C9E2E3F0 F9F3C940 C1F0F4D7 *MSGs.....01..\IST093I A04P*
02:05:56.883 D4E2C7F2 F1F6F140 C1C3E3C9 E5C56B40 D5D6C4C5 40E3E8D7 C5407E40 D7E46DE3 *MSG2161 ACTIVE, NODE TYPE = PU_T*
02:05:57.239 D4E2C7E2 0C000000 00000000 0000F0F1 002DE000 C9E2E3F0 F9F3C940 C1F0F4D3 *MSGs.....01..\IST093I A04L*
02:05:57.239 D4E2C7F2 D5C3C1F3 40C1C3E3 C9E5C56B 40D5D6C4 C540E3E8 D7C5407E 40D3C9D5 *MSG2NCA3 ACTIVE, NODE TYPE = LIN*
02:05:57.306 D4E2C7E2 0C000000 00000000 0000F0F1 0030E000 C9E2E3F0 F9F3C940 C1F0F4D7 *MSGs.....01..\IST093I A04P*
02:05:57.306 D4E2C7F2 E4C3C1F3 40C1C3E3 C9E5C56B 40D5D6C4 C540E3E8 D7C5407E 40D7E46D *MSG2UCA3 ACTIVE, NODE TYPE = PU_*
:

```

Using the timing options

After completing storage analysis, RU count, or VIT extraction, the timing options panel is automatically displayed. Use the timing options to report at certain intervals in the VIT or to process only the VIT records within a certain time range. By default, the entire VIT is processed. All time values, including time stamps, are local (LOC) time. The main panel for storage analysis is shown in Figure 65 on page 395. Select an option and follow the prompts.

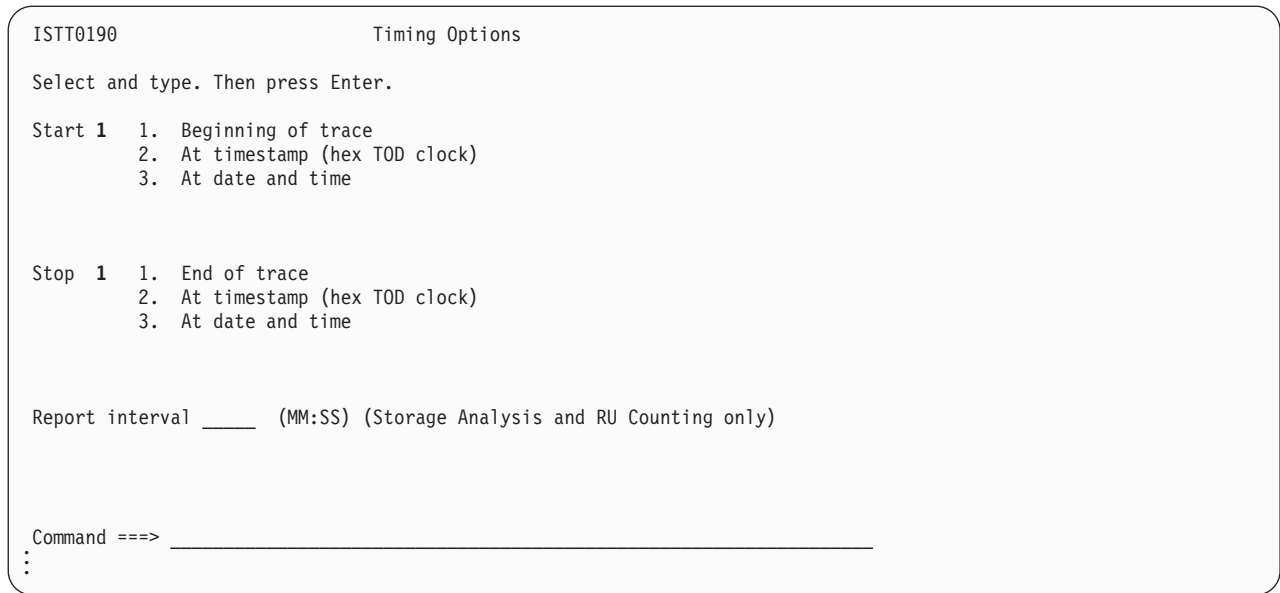


Figure 65. VTAM timing options panel

Note: Report interval, start times, and stop times might not produce the expected output if the VIT was extracted from a dump using the VITAL option. VIT entries in a dump do not have individual time stamps. The VITAL option adds approximated time stamps to the extracted VIT entries based on the time of the dump and times recorded when the VIT wrapped in storage.

Start and stop time

Start and stop times can be at a time stamp within the trace, at a date and time within the trace, or at the beginning of the trace for start and at the end for stop. If you do not specify a start or stop time, the entire VIT is processed.

If you select **At timestamp**, another panel appears on which you may enter a System 370 time-of-day (TOD) time stamp. To ensure that the time stamp reflects your local time, you must add the time zone value to the high-order word. The time zone can be obtained from a dump, if the trace being used has been extracted from it, or by browsing the first time-stamp record from the trace in hexadecimal format. For example:

```

A905470D237491E4   GMT TOD on a GTF trace record
+ FFFFBCF100000000   Time zone (padded with zeros)
-----
A90503FE237491E4   Local TOD

```

You can enter either the high-order 4 bytes of the time stamp, such as X'A90503FE', or all 8 bytes, such as X'A90503FE237491E4'.

Note: All hexadecimal time stamps reported by the VIT analysis tool are local time stamps. (The time zone has already been added.)

If you select **At date and time**, another panel appears on which you can enter both date and time selections.

The time can be:

hh:mm:ss

or

hh:mm:ss.ddd.

The date can be:

Calendar format (mm/dd/yy)

or

Julian format (yy.ddd).

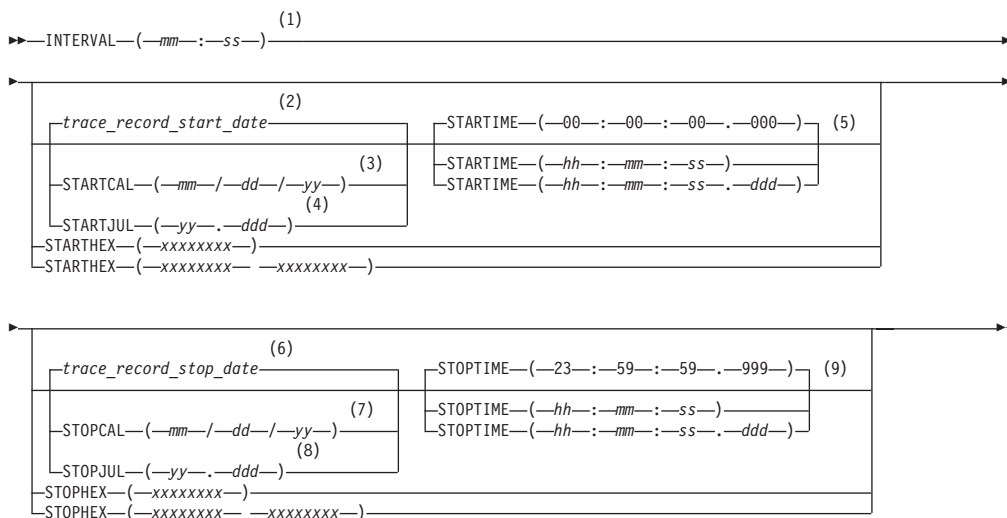
Report interval

If you select report interval, a report is written for each time interval determined from your selection. For example, if you have a storage analysis trace that was recorded for a two-hour period and you specify an interval of five minutes, you get a report of the storage allocated and freed for every 5-minute period of the trace. A total for the whole two hours is printed regardless of the interval specified.

Note: Interval is valid only for storage analysis and RU counting.

Parameter syntax

Use the following syntax only if you are using an editor to write your parameters as an alternative to the panel interface. See "How to read a syntax diagram".



Notes:

- 1 INTERVAL cannot be used with VIT extraction. INTERVAL can be used with any start or stop time specification. Intervals (mm:ss) are interpreted as 00 through 99 for minutes, and 00 through 59 for seconds.
- 2 If you do not code a start date, the VIT analysis tool uses the date on the first trace record.
- 3 The calendar date (mm/dd/yy) is 01 through 12 for month, 01 through 31 for day, and 00 through 99 for year. For dates, yy values 43–99 are interpreted as 1943–1999, and yy values 00–42 are interpreted as 2000–2042.
- 4 The Julian date (yy.ddd) is 00 through 99 for year and 001 through 366 for day. For dates, yy values 43–99 are interpreted as 1943–1999, and yy values 00–42 are interpreted as 2000–2042.

- 5 Calendar time (*hh:mm:ss.ddd*) is 00 through 23 for hour, 00 through 59 for minutes, 00 through 59 for seconds, and 000 through 999 for milliseconds.
- 6 If you do not code a stop date, the VIT analysis tool uses the date on the first trace record.
- 7 The calendar date (*mm/dd/yy*) is 01 through 12 for month, 01 through 31 for day, and 00 through 99 for year. For dates, *yy* values 43–99 are interpreted as 1943–1999, and *yy* values 00–42 are interpreted as 2000–2042.
- 8 The Julian date (*yy.ddd*) is 00 through 99 for year and 001 through 366 for day. For dates, *yy* values 43–99 are interpreted as 1943–1999, and *yy* values 00–42 are interpreted as 2000–2042.
- 9 Calendar time (*hh:mm:ss.ddd*) is 00 through 23 for hour, 00 through 59 for minutes, 00 through 59 for seconds, and 000 through 999 for milliseconds.

Using the I/O options

The I/O Options panel is displayed automatically after the Timing Options panel. Use the I/O options to designate a wrapped trace, whether you want formatted output, or to specify a title or description for the first page of a report. The main panel for the I/O options is shown in Figure 66. Select an option and follow the prompts.

```

ISTT0022                Input/Output Options

Type, then press Enter.

Trace wrapped?  2  1. Trace wrapped
                 2. Trace did not wrap

Format output?  1  1. Format the VIT entries
                 2. Do not format
                 3. Create trace data set

Title . . . . . Trace Analysis

                        Description
_____
_____
_____

Command ==> _____
:

```

Figure 66. VTAM I/O options panel

Trace wrapped

Select *Trace wrapped* if the trace wrapped when it was recorded. Wrapped means that the trace ran out of space on the specified device and began to write over previously recorded data. A trace recorded on a DASD device is large enough to wrap. A trace recorded on a tape will not wrap. A trace created using VTAMMAP VITAL is not wrapped. (VITAL unwraps the internal VIT when it copies the trace records.)

Format output

VIT entries can be written formatted, unformatted or as a hexadecimal string. Here is a GBLK entry in each format:

- **Do not format** (parameter syntax: **NOFORMAT**). Shows the 32-byte VIT entry as eight words in hexadecimal format, followed by the same 32 bytes as the EBCDIC eye-catcher. For example:

```
C7C2D3D2 0C000000 06638480 06357218 823D6340 00000098 864EBCF8 00000098 *GBLK.....d.....b... ..qf+.8...q*
```

- **Format the VIT entries** (parameter syntax **FORMAT**). Shows the 32-byte VIT entry with labels for each field. For example:

```
GBLK      ASID 0C  RC  00      ID  00      VTA  00      AREA 06638480 ANCH 06357218
          ISSR 823D6340      LEN  00000098 INIT  864EBCF8 RLEN  00000098
```

- **Create trace data set** (parameter syntax **TRACEFORMAT**). The 24-byte GTF header followed by the 32-byte entry in hexadecimal format. This output is not suitable for printing, but is usable as input to the VIT analysis tool or other tools. For example,

```
..u..e5et....7..VTAMTST GBLK.....d.....b... ..qf+.8...q
```

Parameter syntax

Use the following syntax only if you are using an editor to write your parameters as an alternative to the panel interface.



Operands

Note:

1. If **FORMAT** is specified and the VTAM format module is not found or is unusable, a message is issued and the job is stopped. To recover, do either of the following actions and rerun the job:
 - Find the current version of AMDUSRFD and add a STEPLIB DD statement to your application.
 - Specify **NOFORMAT** to print the VIT entries in hexadecimal with an eye-catcher.
2. The **TRACEFORMAT** option is valid only for VIT extraction, and results in output in the same format as that of the TRACE input data set.
3. If you specify the **TRACEFORMAT** option, the output from the VIT analysis may be processed by:
 - The VIT analysis tool
 - ACF/TAP
 - IPCS GTFTRACE, unless the VIT was recorded internally and was extracted from a dump by the VTAMMAP VITAL function
4. The 20 characters that follow the word **TITLE** are used as the title. The rest of the input line is ignored.
5. **DESC** may contain up to 75 additional characters, and you can code up to 4 **DESC** parameters. **DESC** must be the only option on the line and must be the first 4 characters on the line followed by at least one blank.

Creating your own parameter data set

The PARM data set must have fixed records (can be blocked) with:

- LRECL=80
- RECFM=FB

The parameters can be coded in any order, and in lowercase, uppercase, or mixed case. Code only *one* function parameter (SMS, RU, or VITEXT) per job or execution. See the parameter syntax for each trace function for a list of the possible parameters that may be coded.

Restriction: The parameters shown in Figure 67 can be coded in parts to avoid exceeding the maximum line length of 80 characters, but a single parameter cannot be continued on the next line.

```
SMS(GBLK(0,1,...255))
SMS(VTAL(0,1,...255))
SMS(REQS(bufname,...))
DESC ccc...
RU(NC(xx,xx,...xx))
RU(SC(xx,xx,...xx))
RU(DFC(xx,xx,...xx))
RU(FMD(xxxxxx,xxxxxx,...xxxxxx))
RU(AMRU(xxxx,...xxxxxxxx))
RU(xx,...xxxxxxxx)
VITEXT Boolean expression
```

Figure 67. Parameters coded on multiple lines

For example, SMS(GBLK(0,1,2,3)) can be split onto multiple lines as follows:

```
SMS(GBLK(0))
SMS(GBLK(1))
SMS(GBLK(2))
SMS(GBLK(3))
```

Note:

1. If you use an editor, you may include comments in your parameter data set. An asterisk in column one identifies a line as a comment line. If you use the panel interface, you cannot enter comment lines.
2. You may use the DEBUG option to gather information to solve problems with the tool itself. To use it, enter DEBUG with one of the VIT analysis options.
3. Lines cannot be continued. Each parameter must be fully specified on one line.

The following information shows a sample parameter data set:

```
SMS(VTAL) sms(match) SMS(Lengths)
INTERVAL(00:15) STARTIME(12:42:14) STOPTIME(14:00:00)
DESC This is an analysis of the VTAL and VTFR VIT entries.
DESC This will tell the high-water mark (the most storage used)
DESC from the start time to the stop time.
DESC This job will also report the storage used in 15-second intervals.
```

See “Analyzing storage” on page 369, “Counting request/response units (RUs)” on page 377, or “Extracting information from the VIT” on page 385 for additional explanation of the parameters.

Also see “Using the timing options” on page 394 and “Using the I/O options” on page 397.

Chapter 9. Using other problem-solving tools

Many different service aids are available to help you collect information about SNA network problems. This information describes when to use the following aids:

- “Alert messages from NCP”
- “Recording NMVT alerts in LOGREC”
- “Messages issued for 3745 bus switching” on page 402
- “Hardware error recording” on page 403
- “Logical unit connection test (IBMTTEST)” on page 403
- “NCP error recording” on page 404
- “Patch areas” on page 404
- “Using save-area module linkage conventions—Subarea” on page 405
- “Using save-area module linkage conventions—APPN” on page 407

Alert messages from NCP

The NCP in a 3720, 3725, or 3745 communication controller sends alert messages to the VTAM program whenever a serious or permanent error occurs in the communication controller.

The NCP sends hardware error records to the maintenance operator subsystem (MOSS). If the MOSS determines that they are permanent errors, it sends them back to the NCP, which forwards them to all owning host processors.

If a communication network management (CNM) application program, such as the NetView program, is active and authorized to receive alert messages, VTAM forwards the alert messages to that program. Otherwise, VTAM sends a message to the operator's console.

Recording NMVT alerts in LOGREC

A network management vector transport (NMVT) is an SNA request unit (RU) that contains solicited or unsolicited data, such as line statistics and generic alerts. LOGREC is a host data set that contains records of various types of system failures, both hardware and software.

VTAM records all *unsolicited* NMVT alerts and *all* NMVT alerts from local area networks in LOGREC as miscellaneous data records (MDRs). If you have the NetView program, VTAM also forwards the NMVT alerts to the NetView hardware monitor for recording. The NetView program interprets the error information for its operator panels. To determine what document contains more information on NetView's presentation of generic alerts, see Table 48 on page 649. For more information on generic alerts generated by First Failure Support Technology (FFST), refer to *z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT*.

VTAM identifies and records an NMVT alert as follows:

- Each NMVT has an SNA network services (NS) header of X'41038D'.
- Each NMVT that is an alert has a management services (MS) major vector of X'0000'.

- To determine the type of device that sent the NMVT alert, VTAM checks the product ID subvector (X'11') of the NMVT for the hardware machine type. Each type of device has its own unique machine type; for example, each NMVT alert that comes from a 3745 has a machine type of C'3745'.
- VTAM records the NMVT alert in LOGREC as an MDR (type=X'91') with a device type of NMVT (X'30').

You can format and print NMVT alerts from LOGREC using the Environmental Record Editing and Printing program (EREP). See Table 48 on page 649 to determine what document contains information on how to use EREP.

Messages issued for 3745 bus switching

A 3745 can be configured with one or two CCUs. For twin-CCU 3745s, NCP V5 supports the two CCUs as independent communication controllers.

A twin-CCU 3745 can be operated such that, if one of the CCUs fails, the maintenance and operator subsystem (MOSS) switches the I/O controller (IOC) buses from the failing CCU to the active CCU. The operator can then activate all or part of the resources of the failing CCU to the other CCU.

The bus switch can also be performed at the operator's request.

If the bus-switch occurs in the middle of a VTAM channel operation, such as a Read or a Write, VTAM issues messages IST881I and IST882I.

Message IST881I

Message IST881I tells the operator that VTAM either was unable to contact a link station, or lost contact to a link station.

This message is issued for one of the following reasons:

- A bus is being switched from one CCU to the other.
- A dump or load is being done on another channel.
- A dump or load is being done by a local disk.

When the link station becomes available (as indicated in other messages), VTAM resumes CONTACT processing.

To terminate CONTACT processing before the link station becomes available, issue a VARY INACT,FORCE command to deactivate the link station.

Message IST882I

Message IST882I tells the operator that VTAM is waiting for a device end from the link station identified in message IST881I.

You should check to see if the link station is online. If it is, then there is a possibility that NCP is being dumped or loaded over another channel adapter, and no further action is necessary.

Note: If the link station is not operating, not physically connected, or not online, VTAM never receives the device end. In those cases, you should issue a VARY INACT,FORCE command to deactivate the link station.

Hardware error recording

During error recovery processing (ERP), VTAM writes outboard recorder (OBR) records and miscellaneous data records (MDRs) to LOGREC. OBR records are written for hardware errors on channel-attached devices. (OBR records are written for communication adapter-attached devices as well.) MDRs or alerts are written for hardware errors on NCP-attached devices. See “Recording NMVT alerts in LOGREC” on page 401 for more information on MDRs.

EREP formats and prints the LOGREC data set.

OBR records contain information about the following items:

- Sense and status data on all channel-attached devices
- Failures on teleprocessing devices
- Temporary or intermittent failures on I/O devices
- End-of-day requests
- Permanent channel and device errors (unrecoverable errors and unit checks)

Permanent error records show the date, time, logical unit name, type of record, contents of counters, failing CCW, channel device name, CSW, sense information, device type, and flags. The time field shows the time at which the permanent error occurred.

Counter overflow and end-of-day records show the date, time, logical unit name, type of record (counter overflow or end-of-day), contents of counters, channel or unit address, and device type. The time field shows the time at which the counter overflow or end-of-day error occurred.

Counter overflow records are written when the temporary error counter or a device statistics table counter is about to overflow. VTAM maintains a counter for each channel-attached device. This counter tracks temporary errors. Counters of unit check errors by error type are also maintained in the device statistics table.

End-of-day records are written whenever a VARY INACT command is entered for a link or channel.

MDRs contain the following information:

- Statistics on the overflow of error counters for communication controllers
- Record maintenance statistics (RECMS) RUs
- Permanent errors on NCP-attached devices

See the EREP entry in Table 48 on page 649 to determine what document describes how to print and interpret MDR and OBR records.

Logical unit connection test (IBMTEST)

You can enter the IBMTEST command from a terminal to find out whether that terminal can communicate with its owning SSCP. When you use the IBMTEST command, an unformatted RU is sent through the network path supporting the LU-SSCP session. This RU contains the IBMTEST command followed by the number of times the SSCP is to return (echo) the data to the logical unit and optional data (up to 247 bytes) being sent to the SSCP.

You can increase the possibility of repeating an intermittent error that is hard to re-create by using IBMTEST, because you can request up to 255 echoes. You can also use it to determine whether a suspended LU-LU session is caused by either a hardware problem or by a problem with VTAM or an application program.

Start this test with the following command:

```
IBMTEST [n][,data]
```

n Specifies the number of times the test data should be returned to the terminal. Specify *n* as a decimal number in the range 1–255. If no value is specified, a value of 10 is used by default.

data

Specifies the test data to be returned. Specify a character string of up to 247 characters, or the maximum message length of the terminal, whichever is smaller. If no test data is supplied, VTAM returns the following alphanumeric sequence:

```
ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789
```

The IBMTEST command is valid only for terminals that use the USS LOGON format. The IBMTEST command must be defined in the USS table for that terminal.

Note: This echo check does not verify that a terminal can establish a session with an application program in the host, because the session request may specify a different network path than the one supporting the terminal's LU-SSCP session. If the requested path is unavailable, the session request is rejected, even though another path is available.

If there are any errors, the VTAM operator receives a message that contains the logical unit name associated with the terminal, the number of echoes that took place before the I/O error, and the error sense code.

NCP error recording

NCP error recording procedures create record maintenance statistics (RECMS) RUs that contain:

- The initial error status that began the recovery process
- The final error status that caused the permanent RECMS RU to be generated

RECMS RUs are created for each adapter check, program check, unresolved interruption, counter overflow, I/O operation, and permanent line error. The RECMS RUs, which contain the miscellaneous data record (MDR), are sent to the VTAM host that owns the failing component. VTAM then records the MDR and writes the error information to the LOGREC file.

Patch areas

Patch areas are available for VTAM and TSO/VTAM.

VTAM patch area

VTAM supplies a fixed patch area as a separate module. You can insert short service programs in this area to do maintenance-related functions. ISTPATCH is loaded into LPALIB during initialization of the operating system and is pointed to

by the ATCPTCHA field in the ATCVT. The initial size of the patch area is 64 bytes, but you can change the size by link-editing a module named ISTPATCH of the required size.

Code the necessary change in the patch area. (ISTPATCH follows the same coding rules as other modules in LPALIB.) Then replace part of the failing code with a branch to the patch area, allowing you to bypass the failing code.

TSO/VTAM patch area

TSO/VTAM maintains a patch area in each module. The size of the patch area varies from module to module. If you need more information on using these patch areas, contact the IBM Support Center.

VTAM load module list

VTAM has a module list pointed to by ATCMDLST in the ATCVT. Each 16-byte entry in the list contains the following information in the form of:

```
XXXXXXXXXXXXXXXXAAAA
```

where:

- XXXXX is the five significant characters of the module name.
- YYYYYYY is the PTF level (or Julian date if PTF level is not present).
- AAAA is the address of the module in storage.

The following information shows an example of some module list entries:

```
C1C9C3C1 D9E4E8F9 F3F7F4F4 00D8ACE8 *AICARUY93744.Q.Y*
C1C9C3E5 C340F9F1 4BF2F0F4 00D8C490 *AICVC 91.204.QD.*
C1C9C3C9 D6E4E8F9 F4F2F8F9 00D8A4C8 *AICIOUY94289.QuH*
C1C9C3C9 D9E4E8F8 F4F2F9F3 00D87000 *AICIRUY84293.Q..*
C1C9C3E7 D440F9F1 4BF0F8F9 00D8C288 *AICXM 91.089.QBh*
C1D7C3D2 E440F9F1 4BF2F9F5 00D885C0 *APCKU 91.295.Qe.*
C1D7C3D9 E440F9F1 4BF0F9F2 00D8C190 *APCRU 91.092.QA.*
C1D7C3E2 D940F9F1 4BF0F8F9 00D8BE60 *APCSR 91.089.Q.-*
C1D7C3E2 E4E4E8F8 F5F3F8F7 00D8B9D8 *APCSUUY85387.Q.Q*
C1D7C3E4 C5E4E8F9 F2F9F1F5 00D88E78 *APCUEUY92915.Q..*
```

You can use this module list table to:

- Determine issuer entries in VIT records
- Search for save-area base registers for modules that reside in LPA
- Verify PTF levels of modules

Using save-area module linkage conventions—Subarea

VTAM traces the flow of the execution of three VTAM components, SSCP, PUS, and LUS, by saving the work areas of modules in these components. The addresses of the module work areas are stored in either of these control blocks:

- Network configuration services parameter list (NCSPL)
- Request/response unit processing element (RUPE)

In the RUPE, the work area address can be found at RUPEDAP. In the NCSPL the work area address can be found at NCSPLWKA. For the hex offsets of these fields, see z/OS Communications Server: SNA Data Areas Volume 1.

The NCSPL or RUPE work area contains the work and save-areas for each module invoked for the command that the NCSPL or RUPE represents. The module work and save-areas provide status information that pertains to both the processing of that command and any interruptions in the processing.

This status information includes a record of which modules were entered, which modules returned to their callers, and which modules returned with a return code. Each module save-area contains the 4th, 5th, 7th, and 8th characters of the module name and the register 15 value that includes a pointer to the last module called by this module. If this address is not in the dump, the module can be obtained by comparing the address to the addresses in the VTAM module list pointed to by ATCMDLST out of the ATCVT. See "VTAM load module list" on page 405.

The high-order byte of the register 15 save-area also indicates the status of the last module called. (In 31-bit mode the address fills register 15, causing the status to overlay the high-order byte of the address in the register 15 save-area.)

Byte value	Status indicated
FE	The called module has returned to this module without a return code.
FF	The called module has not returned to this module.
nn	The called module has returned to this module with a return code of nn.

Figure 68 on page 407 is an example of what the NCSPL or RUPE work area might contain for modules invoked for a VTAM process using save-area module linkage conventions. Using this convention, the save-area contains a 4-byte module identifier, such as ACRT, at the location pointed to by register 13 for each entry in the chain.

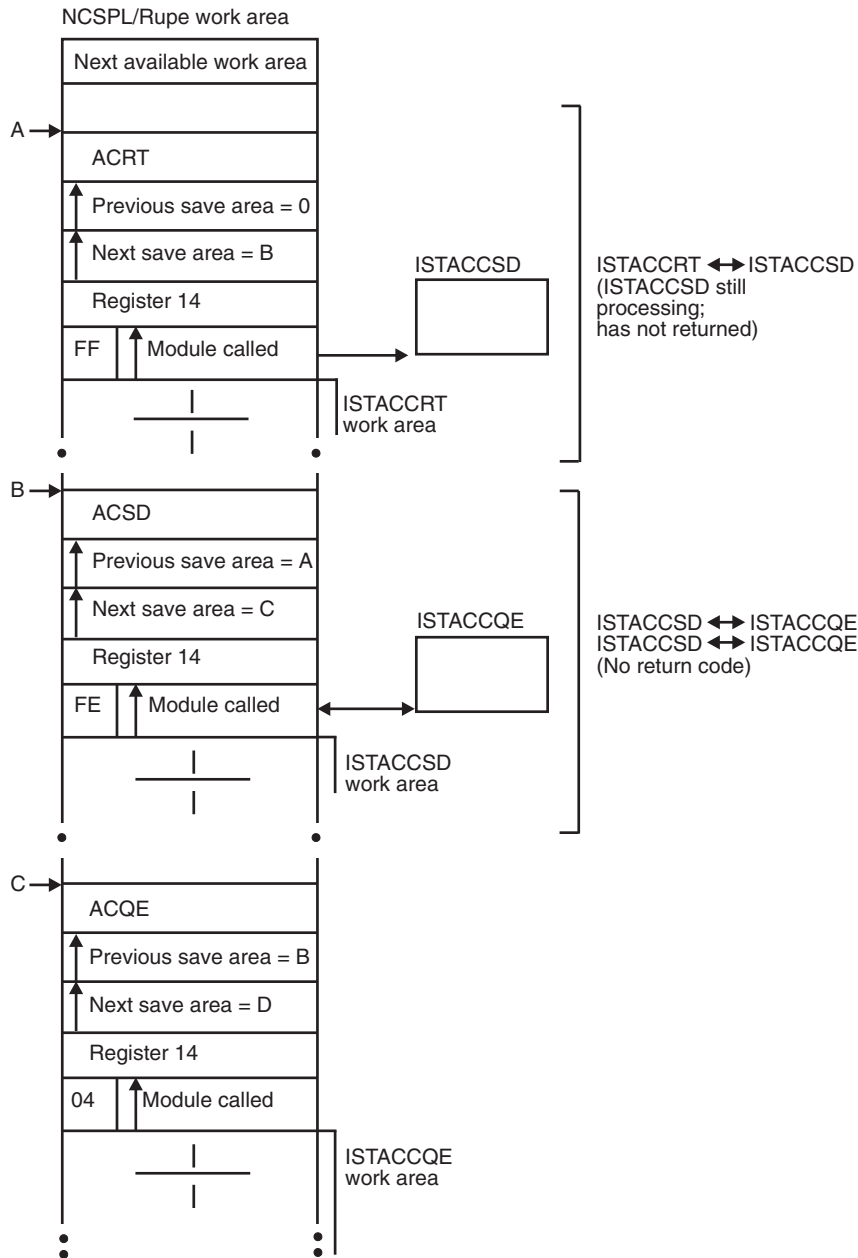


Figure 68. Save-area module linkage conventions—subarea

Using save-area module linkage conventions—APPN

Figure 69 on page 409 shows the save-area module linkage convention for APPN used by some VTAM modules. To determine the convention used for your module, find the location pointed to by register 13 and check the 8-byte field preceding this address. If you find an 8-byte module name such as ISTACCRT, your module was coded using the save-area module linkage convention for APPN. The first three characters will always be IST for a VTAM module.

The addresses of the module work areas for modules using the save-area module linkage convention for APPN are stored in process scheduling services (PSS)

control blocks. The first word of the save-area, pointed to by register 13, contains the pointer to the VTAM work area (VWA) header in the PSS.

For the save-area module linkage convention for APPN, the following save-area format pointed to by the address in register 13 is used:

Offset Contents

X'-08' Module eye-catcher C'XXXXXXXX'

X'+00' Address of ISTVWA

X'+04' Backward save-area pointer (to previous save-area)

X'+08' Pointer to next available area

X'+0C'...
Registers 14 - 12

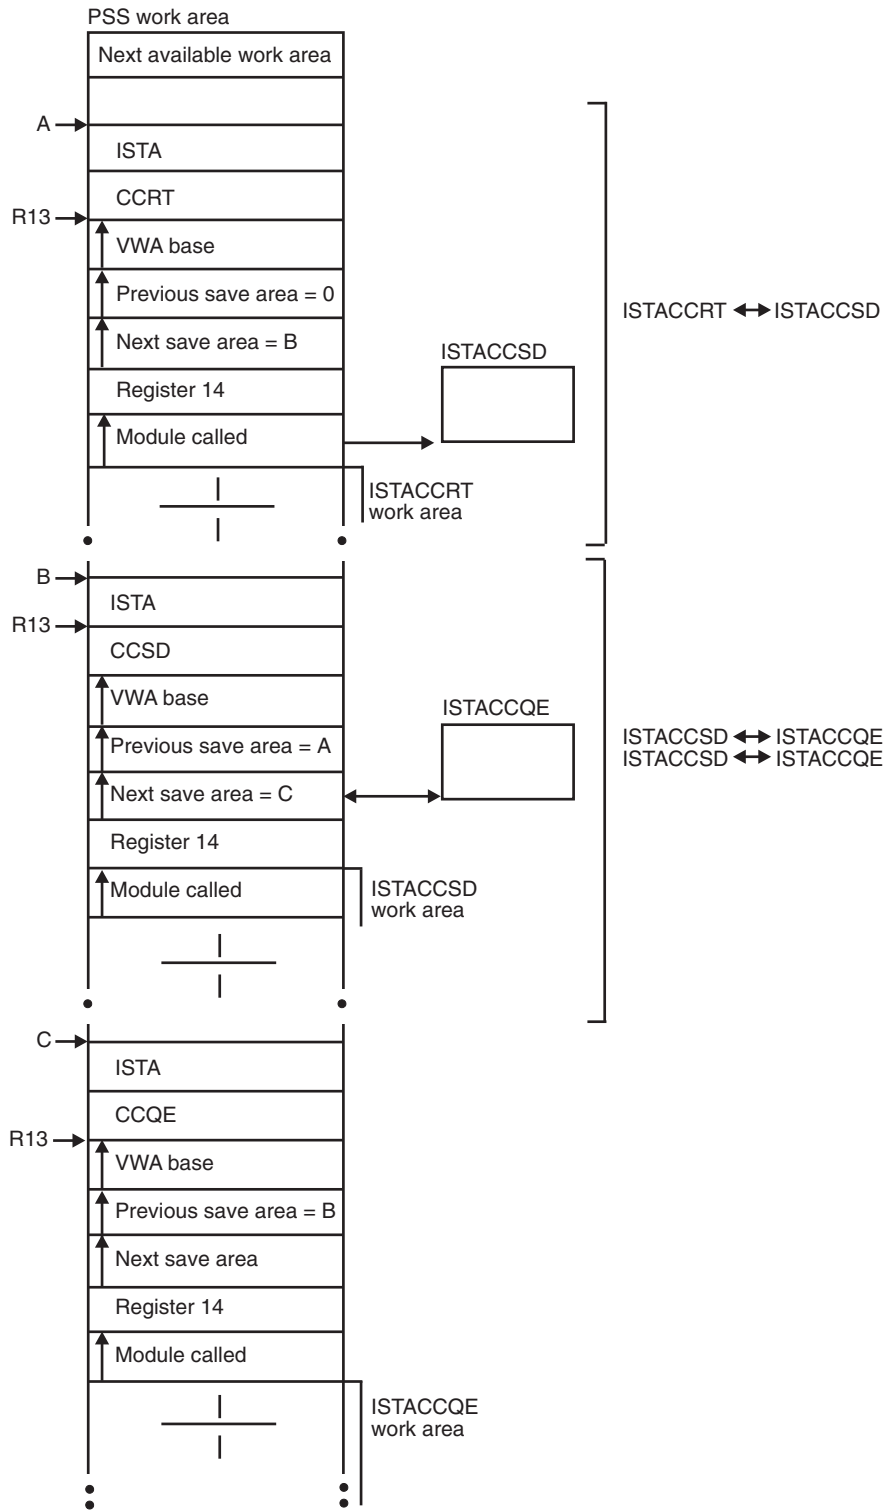


Figure 69. Save-area module linkage conventions—APPN

Part 3. Appendixes

Appendix A. Channel programs

This appendix describes the normal sequence of I/O channel control words (CCWs) within channel programs for the communication controller and channel-attached devices. If you determine that the problem is in an I/O sequence, you need to know the normal sequence of I/O CCWs within these channel programs. With a dump, the VIT trace with the CIO option, or a CCWTRACE (if available), you can compare the sequence that happened with the expected sequence. When there is a deviation, you can then look at status and sense bytes returned from the communication controller or the cluster controller for information that can help you determine the location of the error condition.

This appendix includes the following topics:

- “Channel programs for channel-attached type 2 and type 4 physical units”
- “PUNS-related channel programs” on page 418
- “Channel programs for channel-to-channel adapters (CTCA), multipath channel (MPC), and APPN host-to-host channels” on page 420
- “Channel programs for channel-attached non-SNA 3270 devices” on page 440

Channel programs for channel-attached type 2 and type 4 physical units

The ICNCB represents type 2 and type 4 physical units, and contains addresses and CCWs needed for channel programs. Figure 70 on page 414 shows the following information:

- ICNCB
- Location in storage of various CCWs
- Write buffers required for writing three PIUs, each of which is contained in a single buffer. Write or Write-Break CCWs alternate with transfer-in-channel (TIC) CCWs.

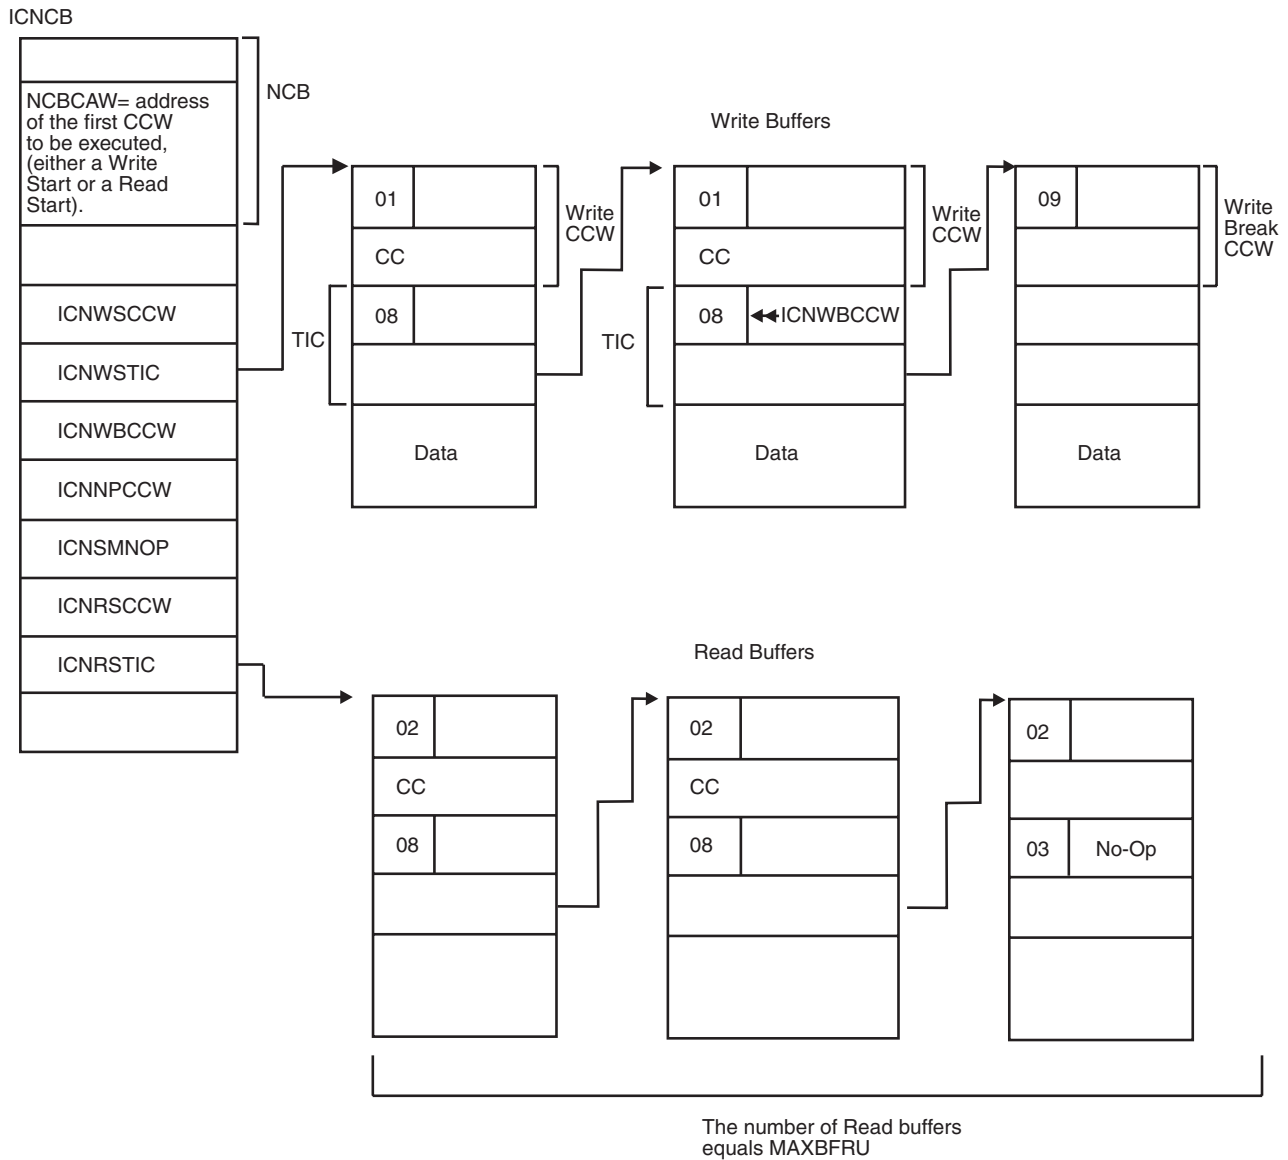


Figure 70. Data areas used by channel programs for PU types 2 and 4

Channel commands for channel-attached type 2 and type 4 physical units

Table 23 describes the channel commands used by the VTAM program to write data to and read data from channel-attached type 2 and type 4 physical units.

Table 23. VTAM channel commands for type 2 and type 4 physical units

Command code (hex)	Command	Description	Contents of address field
01	Write	Transfers data from storage in the host processor to the I/O device.	Output area

Table 23. VTAM channel commands for type 2 and type 4 physical units (continued)

Command code (hex)	Command	Description	Contents of address field
02	Read	Transfers data from the I/O device to storage in the host processor.	Input area
03	No-op	Causes the channel to respond with a channel end and device end. It is the last CCW in a read or write chain.	Zero
04	Sense	Transfers 1 or 2 bytes of sense data to storage in the host processor.	Address of sense data area
08	TIC (Transfer in Channel)	Causes the channel to fetch an instruction that is not the next sequential instruction within the channel program sequence.	Address of next CCW to be executed
09	Write Break	Transfers data from storage in the host processor to the I/O device and indicates that it is the last or only Write command in a chain of Write CCWs.	Output area
31	Write Start 0	Begins a Write sequence. Alternates with Write Start 1.	Zero
32	Read Start 0	Begins a Read sequence. Alternates with Read Start 1.	Zero
51	Write Start 1	Begins a Write sequence. Alternates with Write Start 0.	Zero
52	Read Start 1	Begins a Read sequence. Alternates with Read Start 0.	Zero
Note: Data transfer does not occur on Read-Start or Write-Start commands.			
61	Write XID	The host sends the Write XID command to signal the NCP that a channel contact sequence is beginning and to prepare to receive the host's XID.	Zero
62	Read XID	The host sends the Read XID command to signal the NCP that the host expects to read the NCP's XID.	Zero
A3	Discontact	Indicates that the channel is no longer contacted and the attachment to the transmission group should be broken. Releases the PIUs on the channel hold and intermediate queues.	Zero
C3	Contact	Establishes contact between the host and the NCP. Tells the NCP to use XID information for operations with the host.	Zero
93	Restart	Causes the controller to reset its switches to indicate that the last Write-Start and Read-Start commands were Write-Start-1 and Read-Start-1 commands.	Zero

Format of transfer-in-channel (TIC) CCWs

The Format 1 TIC CCW is formatted as follows:

Byte (hex)

Contents

00 X'08' (TIC identifier)

01-03 Zero

04-07 Real address

A doubleword TIC extension immediately follows both the Format 0 TIC and the Format 1 TIC. VTAM uses the last 4 bytes of the TIC extension to contain the virtual address of the next buffer in the chain.

The TIC extension is formatted as follows:

Byte (hex)

Contents

08-0B Reserved

0C-0F Virtual address

For write buffers, the next to the last physical buffer is handled specially. The real address of the TIC points to the last Write-Break CCW (ICNWBCCW), but the virtual address points to the last write buffer that contains data (see Figure 71 on page 417.) This last buffer is formatted with a Write-Break CCW that is not used but is copied into ICNWBCCW. If only one buffer exists in the channel program, the Write-Start TIC is formatted so that the real address points to ICNWBCCW, but the virtual address points to the only write buffer.

The last write buffer looks unusual because the Write-Break command is chained, but the next CCW is zero. This Write-Break CCW is never physically executed by the channel, but the copied version of the CCW (in the ICNCB) is executed.

Figure 71 on page 417 shows the write buffers required for writing two PIUs when each spans three buffers.

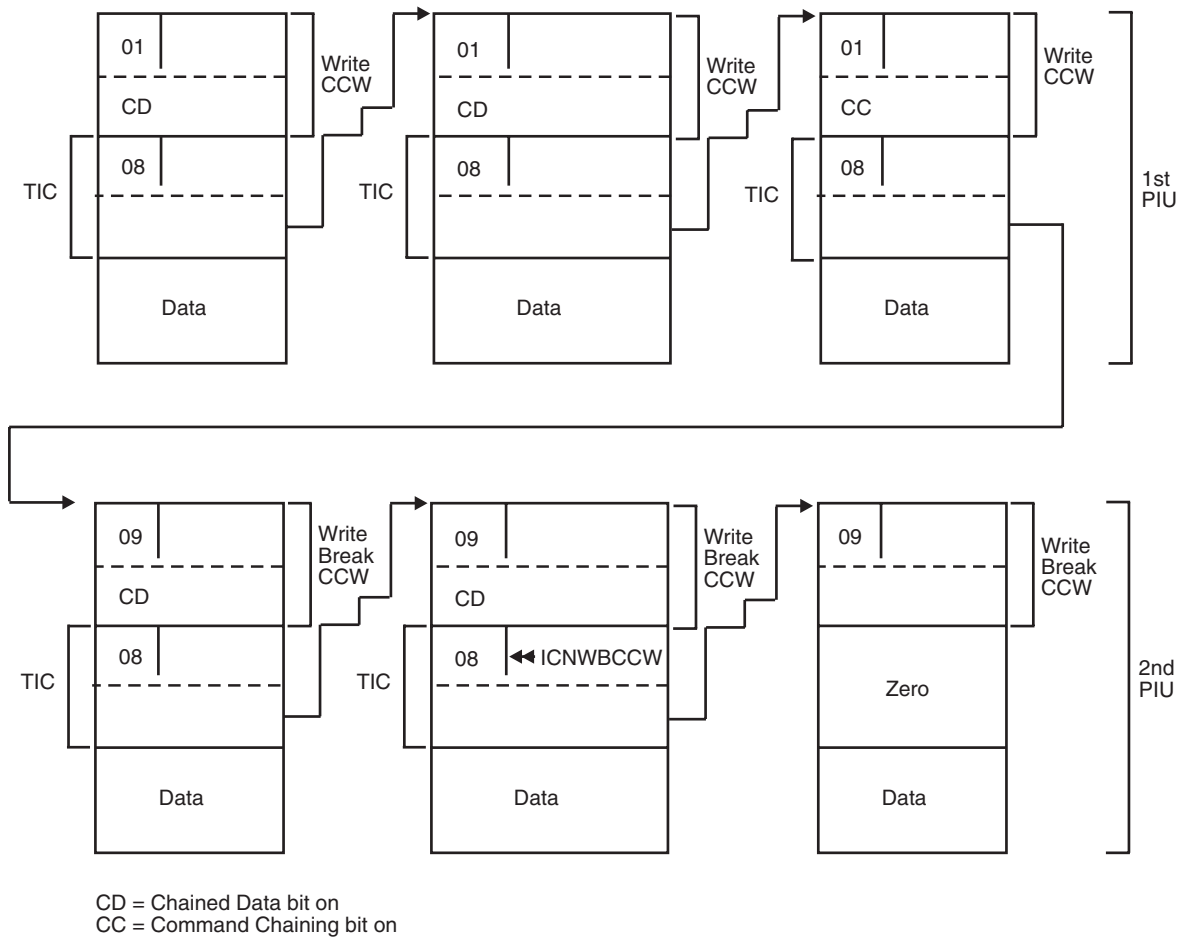


Figure 71. Format of Write CCWs with chained data

Channel program sequences

The following information describes the CCWs, in the order that they are executed, in a typical WRITE channel program.

1. *ICNWSCCW - Write Start*: Each time a Write sequence (the Write start and all associated write buffers) has completed successfully, the Write-Start CCW is alternated: The first Write-Start operation code is X'31' (Write Start 0), the second Write-Start code is X'51' (Write Start 1), the third Write-Start code is X'31', and so on. There is no data transfer associated with this CCW; it is used to inform the controller that the previous Write has successfully completed.
2. *ICNWSTIC - Write Start TIC*: This CCW is used to chain to the first Write CCW. When the WRITE channel program consists of a single write buffer, the real address points to ICNWBCCW (the last Write Break). The virtual address of this CCW always points to the first write buffer on a chain or is zero when no WRITE channel program is pending or active for the device.
3. *Write or Write-Break CCWs Alternating with TICs*: Figure 70 on page 414 shows the write buffers required for writing three Write PIUs, each of which is contained in a single buffer. Figure 71 shows the write buffers required for writing two PIUs, when each spans three buffers.

4. *ICNWBCCW CCW(3) - Last Write Break*: In a channel program this is the last Write CCW that is physically executed by the device. The data address points to the last write buffer that contains data. If a Read is requested, when this operation completes, the controller will signal a status modifier (in addition to channel end and device end). This causes ICNNPCCW to be skipped by the channel.
5. *ICNNPCCW - No-Op*: If a Write completes and no Read is requested, No-Op is the ending CCW in a channel program. If a Read was requested, this CCW is skipped by the channel. The command chain flag will be on in this CCW if a Read was previously requested but was not previously initiated, and read buffers are available.
6. *ICNSMNOP - Status Modifier No-Op*: This CCW receives control after the last Write Break if a Read is signaled. This CCW normally command chains to the Read Start so that writing and reading occur without interruption. If no buffers are available, however, the channel program ends here, and VTAM recognizes that a Read is required.
7. *ICNRSCCW - Read Start*: Each time a Read sequence_ (the Read Start and at least one Read) has completed successfully, the Read-Start CCW is alternated. The first Read-Start generation code is X'32' (Read Start 0), the second Read-Start generation code is X'52' (Read Start 1), the third Read-Start generation code is X'32', and so on.

As with the Write-Start CCW, there is no data transfer. The CCW alternation is used to inform the controller that the previous Read has successfully completed.

When a Write is not required, but a Read has been requested by an attention status, Read Start is the first CCW in the channel program.

8. *ICNRSTIC - Read Start TIC*: This CCW is used to chain to the first Read CCW. The virtual address of this CCW points to the first read buffer on the chain, except:
 - When deblocking PIUs, in which case the virtual address is changed by the channel end appendage
 - When there are not enough read buffers available, in which case it is zero
9. *Read CCWs Alternating with TICs*: Figure 70 on page 414 shows the buffers required for reading three PIUs.
10. *Read No-Op*: This CCW should never be executed. If it is, it indicates that the controller and VTAM do not agree on how many read buffers are required. The normal ending status for a Read is channel end, device end, attention, or unit exception. The unit exception indication is presented on the Read CCW that has completed data transfer and terminates the command chaining. The attention indication is the same as the unit exception indication, but it also means that a Read is requested.

Unit exception may also be presented to the Read-Start CCW. It is used by the controller to release input buffers. If the same Read Start (as opposed to the alternate Read Start) is given to the controller, the data buffers must be resent.

PUNS-related channel programs

During activation of an NCP in a channel-attached communication controller, the SSCP sends a Contact RU to PUNS. PUNS responds by giving control to ISTTSCP4 to schedule one of five channel programs (A–E in Table 24 on page 419). When the SSCP sends a Disconnect RU to PUNS, one of two channel programs (F and G in Table 24 on page 419) is executed.

Table 24. PUNS-related channel programs

	CCW	Code (hex)	Flags	Notes
A.	Sense	04	SLI	Determines if the device needs to be loaded.
B.	Write XID	61	SLI,CC	This channel program follows A if the device does not need to be loaded. It is followed by either C or D.
	No-Op	03	SLI	
C.	Write Break	09	SLI,CC	This channel program follows B if there is no command reject.
	Read XID	62	SLI,CC	
	Read	02	SLI,CC	
	No-Op	03	SLI	
D.	Restart/Reset	93	SLI,CC	Executed only if the Write XID in B caused a command reject (implies the NCP is NCP Release 2 or earlier).
	No-Op	03	SLI	
E.	Contact	C3	SLI,CC	Restart/Reset is executed only if a Contacted (error) response is sent to PUNS.
	Restart/Reset	93	SLI,CC	
	No-Op	03	SLI	
F.	Discontact	A3	SLI,CC	This channel program is executed only if B did not cause a command reject.
	No-Op	03	SLI	
G.	No-Op	03	SLI	This channel program is executed if B caused a command reject.

These CCWs are contained in I/O buffers that are allocated from the IOBUF buffer pool when doing PUNS I/O.

The data area pointed to by the address portion of a Read XID or Write XID CCW is described in ISTXID in z/OS Communications Server: SNA Data Areas Volume 1.

Channel programs for channel-to-channel adapters (CTCA), multipath channel (MPC), and APPN host-to-host channels

The VTAM CTC function supports two protocols: blocking and nonblocking. If VTAM is communicating with another VTAM through the CTC adapter and both VTAMs support the blocking protocol, blocking is the chosen protocol. If one VTAM does not support blocking protocols, the nonblocking protocol is used.

You cannot specify the protocol choice during system definition. Because the blocking protocol is the preferred mode, it is used if both VTAMs support it.

Channel programs for activating the CTC connection

Each side of a channel-to-channel adapter (CTCA) is represented by a cross-channel node control block (XCNCB). In addition, each side has a physical unit service I/O (PIO) control block and a station control block (SCB). The PIO is used for exchange ID (XID) channel programs (the physical unit services I/O that occurs before the link is active). The PIO is mapped by ISTPIO. The SCB is a station work area where CCWs for normal data transfer are built. It is not mapped, and is *not* described in z/OS Communications Server: SNA Data Areas Volume 1 or z/OS Communications Server: SNA Data Areas Volume 2.

A series of channel programs are issued when the operator activates the CTC connection. These I/O exchanges are used by the hosts to communicate various capabilities to the other host. The capabilities are transferred via the XID channel program.

The three primary pieces of information gained through the XID exchange are:

- Choice of protocol – blocking or nonblocking
- Determination of who is X-side and who is Y-side
- I/O buffering information to use when the connection is active

Protocol choice

As mentioned earlier, VTAM always chooses the blocking protocol if the partner VTAM can support it. If the XID indicates that it cannot support, the nonblocking protocol is used.

X-side / Y-side

Figure 72 on page 421 shows that how VTAM determines which side will be the X-side and which side will be the Y-side. In this example, the operator in subarea 4 is the first to activate the link. Subarea 4 begins as the X-side and then switches to the Y-side.

I/O buffering

During the XID exchange, each host informs the other about its read buffer capability.

In the blocking protocol, the total size of the single read buffer is communicated. The write buffer in the other host is allocated based on the size of the Read buffer.

In the nonblocking protocol, the total number of read buffers available, as well as the size of each buffer, is communicated. Each host must then allocate the write portion of its channel program to match the read portion of the other host.

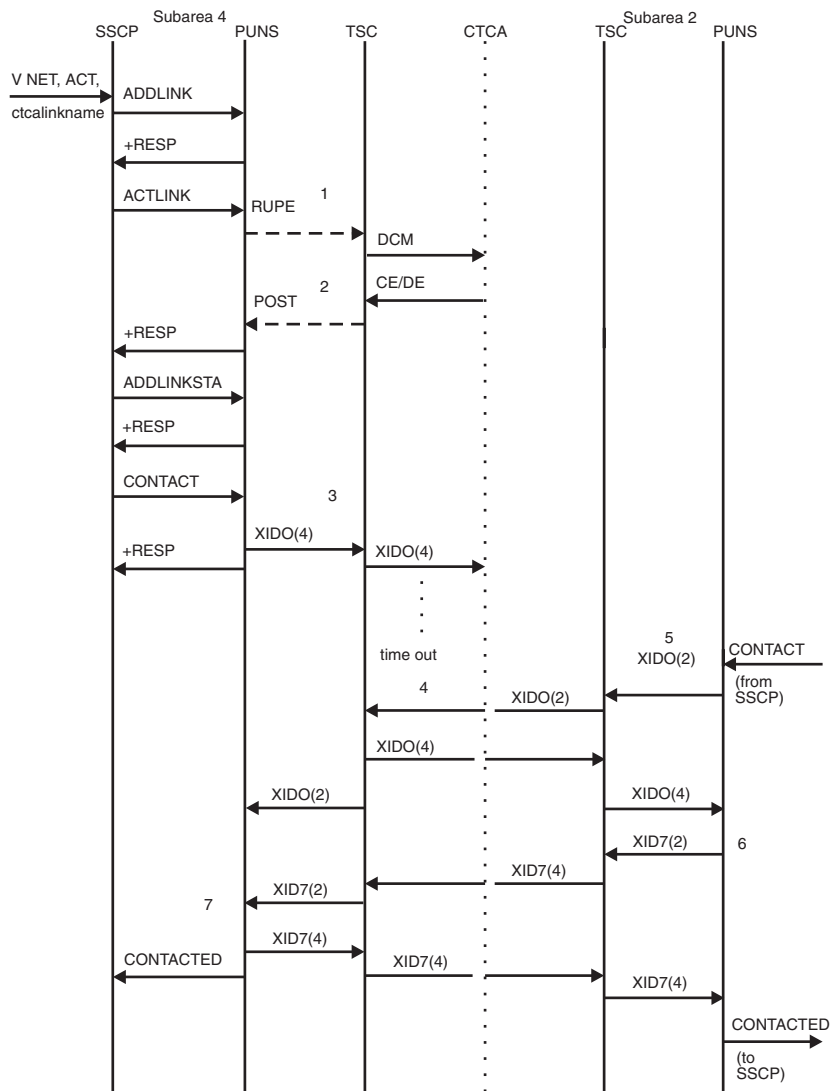


Figure 72. Example of an XID exchange

CTCA Channel-to-channel adapter

PUNS Physical unit services

SSCP System services control point

TSC Transmission subsystem component

Note: (2) and (4) refer to Subarea 2 and Subarea 4.

1 PUNS queues a RUPE to NCBPUPAB, causing TSC to get control.

2 TSC posts PUNS.

- 3 Subarea 4 assumes X-side protocols because it is initiating the XID exchange.
- 4 The XID exchange times out because the link in subarea 2 has not been activated. The CONTACT is issued in subarea 2, and upon attention, subarea 4 assumes Y-side protocols.
- 5 Assume the operator activates the link in subarea 2. Subarea 2 assumes X-side protocols.
- 6 The subarea with the lower subarea number (subarea 2) responds first to the XID0 exchange by sending an XID7. (XID7 is an XID Format 2 with the contact option field set to 7.)
- 7 Subarea 4 checks the XID7 from subarea 2 and responds with its XID7. It also sends a CONTACTED RU to the SSCP.

Channel commands for channel-to-channel (CTC) adapters

Table 25 contains the channel commands used for channel programs between two hosts connected by a channel-to-channel adapter.

Note: The Multipath channel (MPC) connection uses these commands only during activation and deactivation. For the channel commands used for data flow, when the MPC is in the CONTACTED-ACTIVE phase, see “Channel programs for multipath channel (MPC)” on page 434.

Table 25. Channel commands for channel-to-channel adapters

Command code (hex)	Command	Description
01	Write	<p>Transfers data from storage in this host processor to the CTC adapter (CTCA). For CTCAs, this CCW is used to transfer status information from XCNOCTL, XID information from PIOPOTXT, validity checking information from XCNOVTXT, and data from buffers.</p> <p>Status information, mapped by XCNOCTL, is transferred in the first 8 bytes of the write buffer, and only one write buffer is used. Figure 75 on page 427 illustrates the buffer usage.</p> <p>For MPC connections this CCW, preceded by a Prepare command, is used to transfer data from the CPNCB transmit buffer. The transmit buffer is defined by YCNOBUF of the YCNCB contained within the CPNCB. The first 8 bytes of the transmit buffer is mapped by ISTBKHDR and contains control information about the current data transfer. Figure 81 on page 437 illustrates the MPC buffers used for normal data transfer.</p>

Table 25. Channel commands for channel-to-channel adapters (continued)

Command code (hex)	Command	Description
02	Read	<p>Transfers data from the CTCA to this host processor. For CTCAs, this CCW is used to read status information into XCNICTL, XID information into PIOPITXT, validity checking information into XCNIVTXT, and data into buffers.</p> <p>Status information, mapped by XCNICTL, is transferred in the first 8 bytes of the read buffer, and only one read buffer is used. Figure 75 on page 427 illustrates the buffer usage.</p> <p>For MPC, this CCW is used to read normal data into the CPNCB transmit buffer. The transmit buffer is defined by YCNIBUF of the YCNCB contained in the CPNCB. Figure 81 on page 437 illustrates the MPC buffers used for normal data transfer.</p>
03	No-Op	<p>Causes the channel to respond with a channel end or device end. It is the last CCW in a read or write chain.</p> <p>The No-Op command does not apply to the MPC connection.</p>
08	TIC	<p>Causes the channel program to execute an instruction that is not the next sequential instruction within the channel program sequence.</p> <p>The TIC command does not apply to the MPC connection.</p>
14	Sense Command Byte (SCB)	<p>The SCB is normally issued in response to an attention generated when the adapter processes a WCTL from the other side. The SCB clears WCTL from the adapter, allowing the WCTL CCW to complete.</p>
17	Write Control (WCTL)	<p>Causes an attention interruption on the other side of the channel-to-channel adapter. The WCTL is issued to alert the other side that a channel program is active at the adapter.</p>
43	Enable Compatibility Mode (ECM)	<p>Prepares the adapter to operate in System/360 (compatibility) mode. The ECM is issued when VTAM gives up control of the adapter.</p>
C3	Disable Compatibility Mode (DCM)	<p>Prepares the adapter to operate in System/370 (extended) mode. The DCM is issued when VTAM acquires control of the adapter.</p>

Table 25. Channel commands for channel-to-channel adapters (continued)

Command code (hex)	Command	Description
E3	Prepare	Primes the CTC adapter for the next CCW. This CCW does not cause an attention on the other side of the adapter. The Prepare command applies only to the MPC connection.

XID channel program (X-side)

Figure 73 on page 425 shows the data areas associated with the following XID channel program.

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
1	Write Control	17	CC,SLI	Zero	1
2	TIC	08	—	WRITE CCW	—
3	Write	01	CC,SLI	XCNOCTL	8
4	Write	01	CC,SLI	PIOPOTXT	Length of XID
5	TIC	08	—	READ CCW	—
6	Read	02	CC,SLI	XCNICTL	8
7	Read	02	CC,SLI	PIOPITXT	Length of XID
8	TIC	08	—	READ CCW	—
9	Read	02	CC,SLI	XCNVOTXT	4

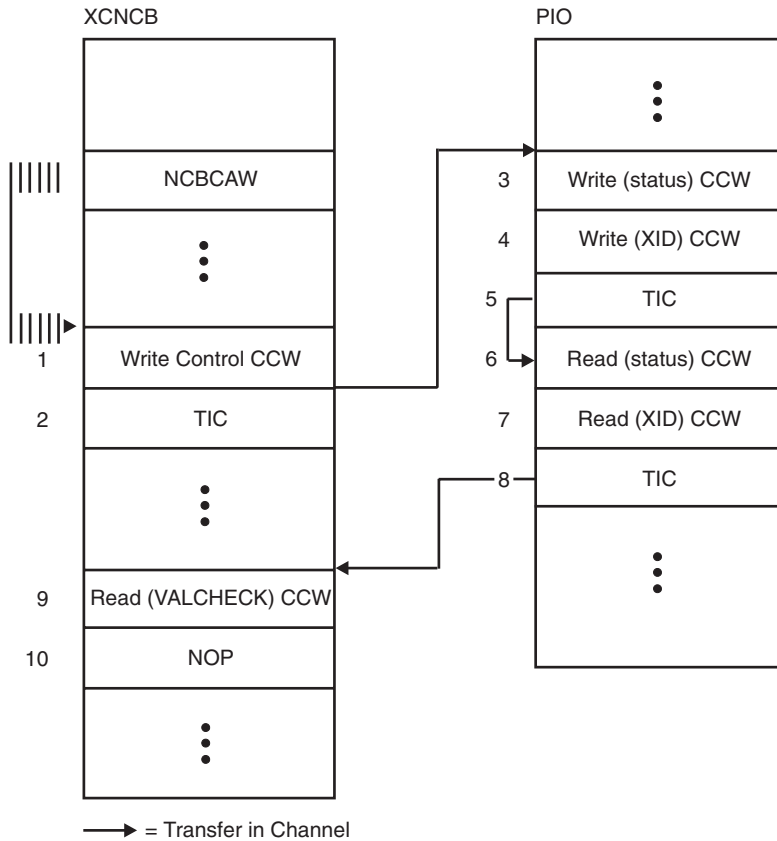


Figure 73. Data areas used for XID channel programs (X-side)

XID channel program (Y-side)

Figure 74 on page 426 shows the data areas associated with the following XID channel program.

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
1	Sense Command Byte	14	CC,SLI	Zero	1
2	TIC	08	—	READ CCW	—
3	Read	02	CC,SLI	XCNICTL	8
4	Read	02	CC,SLI	PIOPITXT	Length of XID
5	TIC	08	—	WRITE CCW	—
6	Write	01	CC,SLI	XCNOCTL	8
7	Write	01	CC,SLI	PIOPOTXT	Length of XID

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
8	TIC	08	—	WRITE CCW	—
9	Write	01	CC,SLI	SCNOVTXT	4

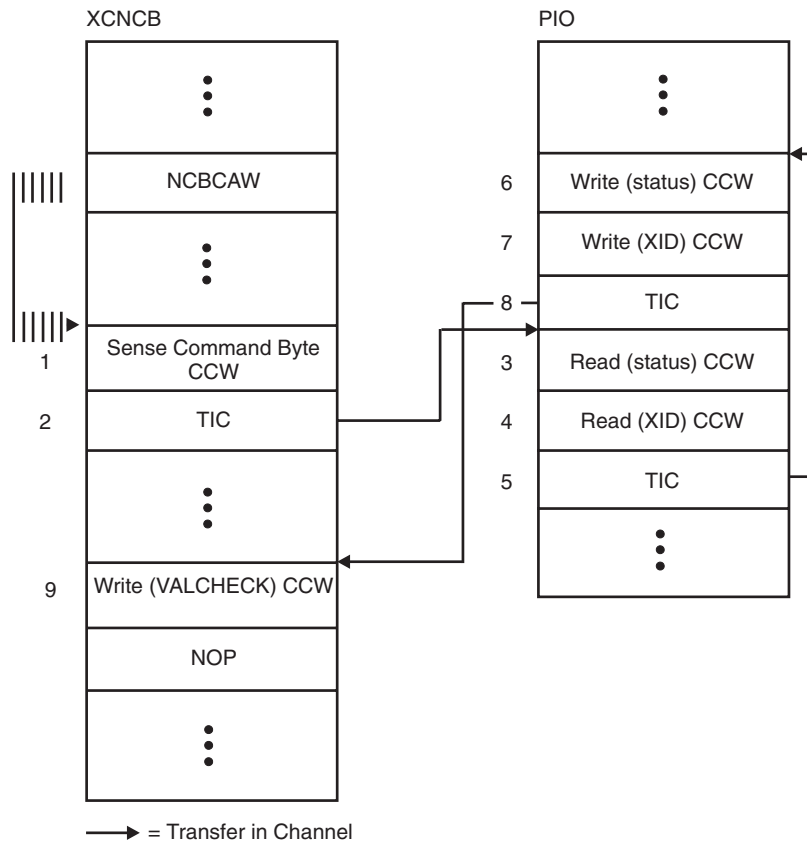


Figure 74. Data areas used for XID channel programs (Y-side)

Channel programs for CTC data transfer: Blocking protocol

MAXBFRU defines a single read buffer. The write buffer in the other host is allocated to be the same size. PIUs are blocked for transfer, and are written by a single write, and read by a single read.

VTAM uses only three CCWs in its Data Transfer channel program: a Write CCW, a Read CCW, and another that is either a WCTL CCW or an SCB CCW. Output control information is included within the write buffer and input control information is received in the first bytes of the read buffer.

Normal data transfer for channel programs is a Write CCW followed by a Read CCW (on the X-side) or a Read CCW followed by a Write CCW (on the Y-side). VTAM uses the procedure described in "Channel programs for activating the CTC connection" on page 420 to determine which side will be X and which will be Y.

During the XID exchange, each host informs the other host of the size (in pages) of the buffer that will be used in the read portion of its channel programs. Each host then allocates a write buffer to match exactly the read buffer of the other host.

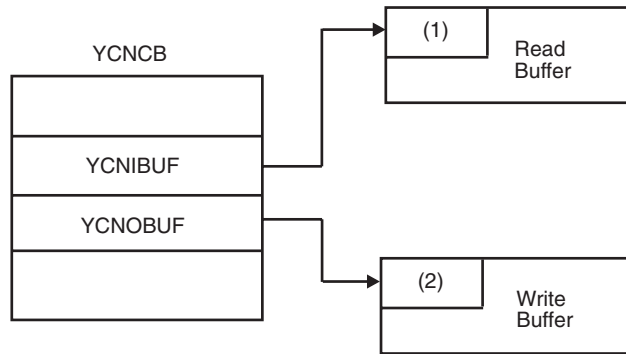


Figure 75. Buffers used for normal data transfer

The numbers 1 and 2 in Figure 75 represent:

- (1) The first 8 bytes of the read buffer is the control information mapped by XCNICTL.
- (2) The first 8 bytes of the write buffer is the control information mapped by XCNOCTL.

Normal data transfer (X-side)

Figure 76 on page 428 shows the data areas associated with the X-side of a normal data transfer channel program.

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
1	WCTL	17/ 14	CC,SLI	Zero	1
2	Write	01	CC,SLI	YCNWRIDA	Number of bytes to transfer

Note: WCTL is used when this host is initiating a write operation. SCB is used when this host is responding to an attention (because the other host has data that it wants this host to read).

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
3	Read	02	SLI	YCNRDIDA	Total length of Read buffer

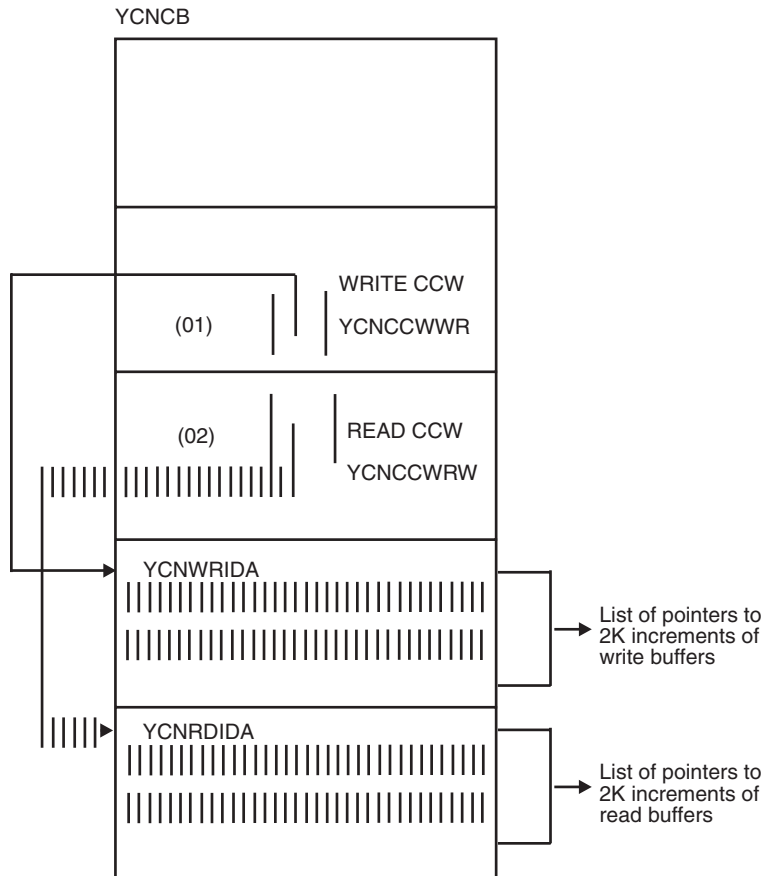


Figure 76. Data areas used for normal data transfer (X-side)

Normal data transfer (Y-side)

Figure 77 on page 429 shows the data areas associated with the following normal data transfer channel program.

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
1	WCTL	17/ 14	CC,SLI	Zero	1

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
Note: WCTL is used when this host is initiating a write operation. SCB is used when this host is responding to an attention (because the other host has data that it wants this host to read).					
2	Read	02	CC,SLI	YCNRDIDA	Total length of Read buffer
3	Write	01	SLI	YCNWRIDA	Number of bytes to transfer

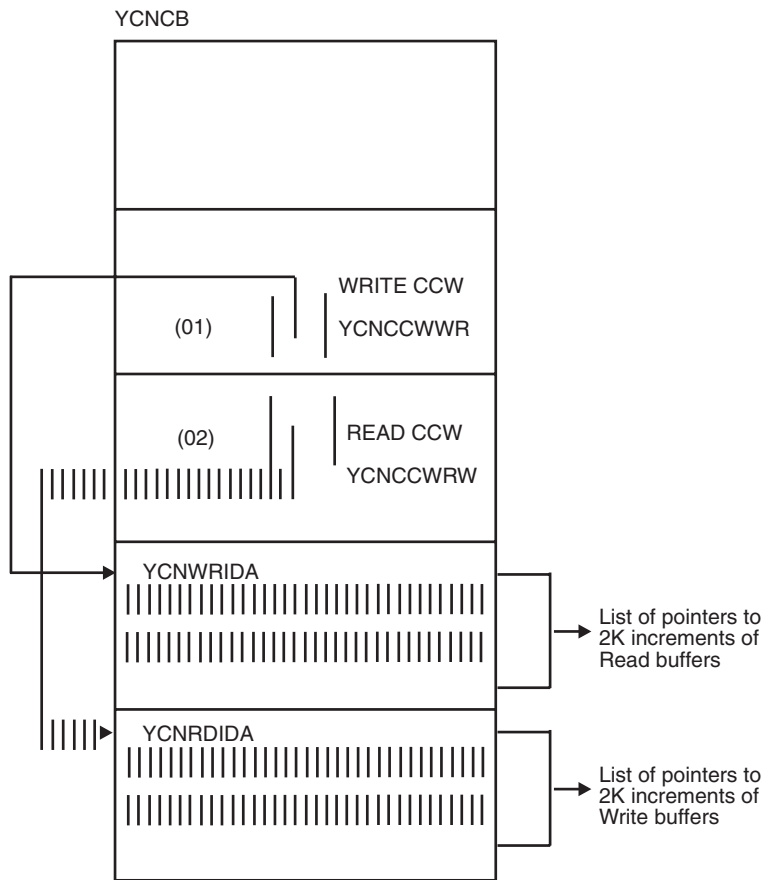


Figure 77. Data areas used for normal data transfer (Y-side)

Channel programs for channel-to-channel adapters: Nonblocking protocol

MAXBFRU defines the total number of IOBUF sized read buffers to use to receive data from the CTC. Consequently, the channel program is a series of writes followed by a series of reads for the X-side. For the Y-side, it is a series of reads followed by a series of writes.

Normal data transfer (X-side) for nonblocking protocols

Figure 78 on page 431 shows the data areas associated with the X-side of a normal data transfer channel program for nonblocking protocols.

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
1	WCTL	17/ 14	CC,SLI	Zero	1
Note: WCTL is used when this host is initiating a write operation. SCB is used when this host is responding to an attention (because the other host has data that it wants this host to read).					
2	TIC	08	—	WRITE CCW	—
3	Write	01	CC,SLI	XCNOCTL	8
4	TIC	08	—	WRITE CCW	—
	Write : :	01	CC,SLI (CD)	Buffer	Length of data in this buffer
5	Write : : Write	01 01	CC,SLI, (CD)	Buffer	Length of data in this buffer Length of data in this buffer
6	TIC	08	—	READ CCW	—
7	Read	02	CC,SLI	XCNICTL	8
8	TIC	08	—	READ CCW	—
	Read	02	CC,SLI	Address of Data	Length of data in this buffer
9	TIC : : Read	08 02	— CC,SLI	READ CCW Address of Data	— Length of data in this buffer
	TIC	08	—	WRITE CCW	—
10	Write	01	CC,SLI	XCNOVTEXT	4

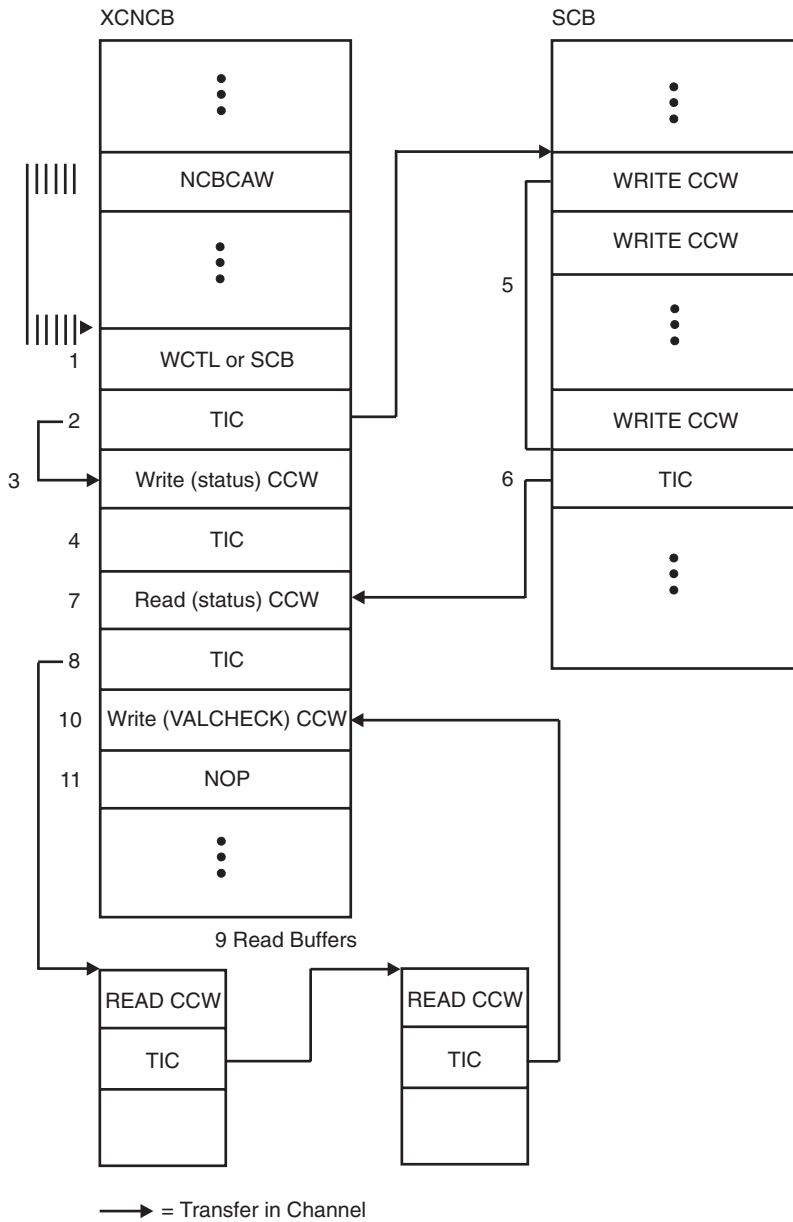


Figure 78. Data areas used for normal data transfer (X-side) nonblocking

Normal data transfer (Y-side) for nonblocking protocols

Figure 79 on page 433 shows the data areas associated with the Y-side of the following normal data transfer channel program for nonblocking protocols.

Sequence	CCW	Command code (hex)	Flags	Address	Byte count
1	WCTL	14/ 17	CC,SLI	Zero	1
Note: WCTL is used when this host is initiating a write operation. SCB is used when this host is responding to an attention (because the other host has data that it wants this host to read).					
2	TIC	08	—	READ CCW	—
3	Read	02	CC,SLI	XCNICTL	8
4	TIC	08	—	READ CCW	—
	Read	02	CC,SLI	Address of Data	Length of data in this buffer
	TIC	08	—	READ CCW	—
5	:				
	Read	02	CC,SLI	Address of Data	Length of data in this buffer
	TIC	08	—	WRITE CCW	—
6	Write	01	CC,SLI	XCNOCTL	8
7	TIC	08	—	WRITE CCW	—
	Write	01	CC,SLI, (CD)	Buffer	Length of data in this buffer
8	Write	01	CC,SLI, (CD)	Buffer	Length of data in this buffer
	Write	01	CC,SLI	Buffer	Length of data in this buffer
9	TIC	08	—	READ CCW	—
10	Read	02	CC,SLI	XCNIVTXT	4

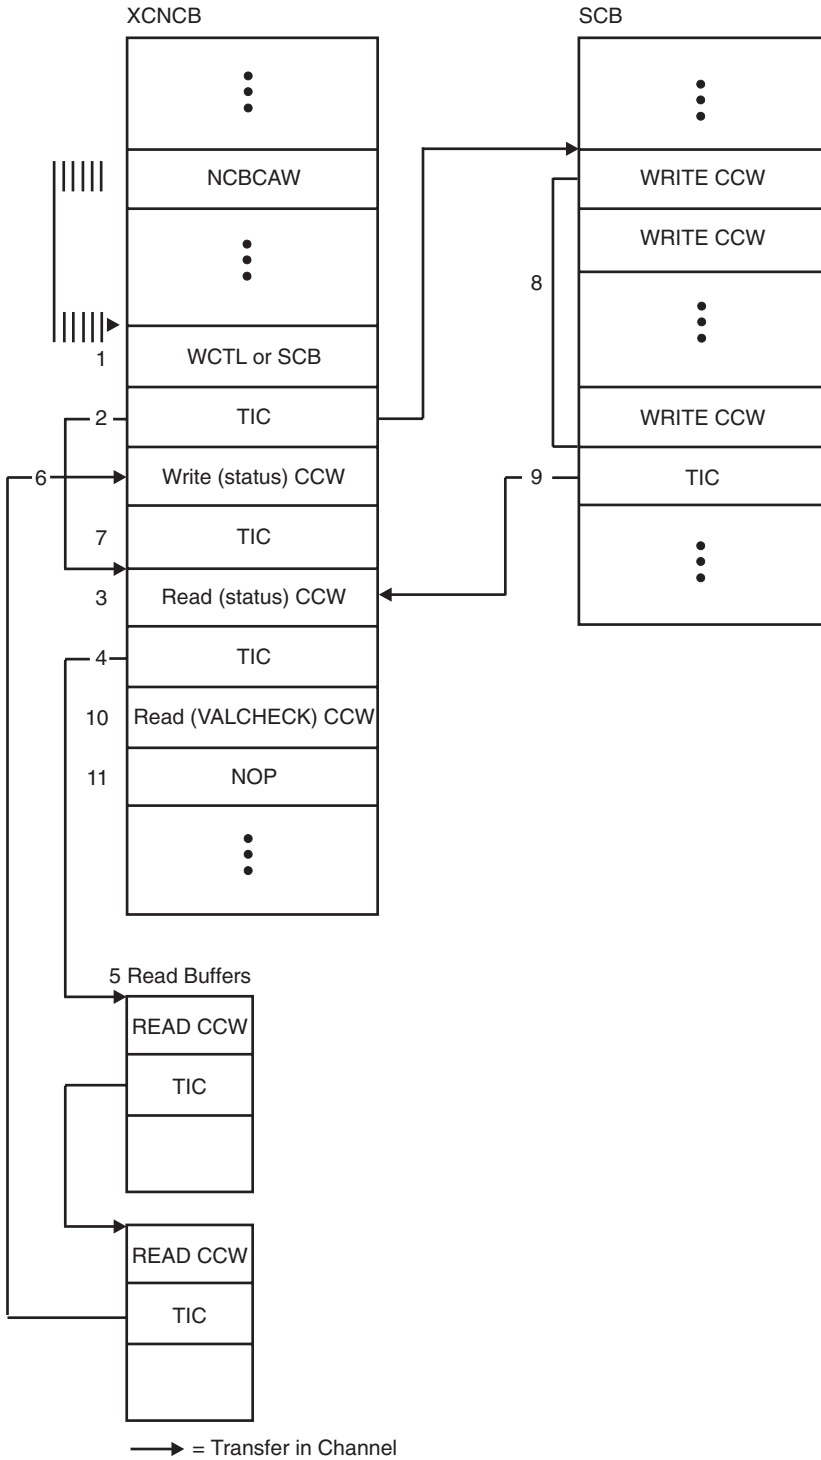


Figure 79. Data areas used for normal data transfer (Y-side) nonblocking

Channel programs for multipath channel (MPC)

The VTAM MPC function is derived from the VTAM channel-to-channel (CTC) function. Both MPC and CTC communicate with the CTC Adapter, but MPC uses its own set of channel programs.

During normal data transfer, MPC uses the never-ending Read channel program. This means that a subchannel defined as a Read device can have an outstanding Read channel program or can be processing the last channel program to complete. On subchannels defined as Write devices, a Write channel program is generated when data is available to be written to the CTC adapter.

For HPDT, a seldom-ending channel program scheme is implemented. For a read subchannel, there are at least four (and possibly up to eight) read CCWs in the channel program, command-chained together. For a write device, there may be as few as one write in the channel program, but there may be up to seven writes command-chained together, depending on the amount of outbound traffic. Program-controlled interrupt (PCI) is used for both the read and write subchannels. Suspend/Resume is used only for write subchannels.

The information that follows describes the multipath channel programs for activating or deactivating an MPC connection as well as for normal data transfer.

Channel programs for activating the MPC connection

A series of channel programs is issued when the operator activates the MPC connection. These I/O exchanges are used by the hosts to communicate various capabilities to the other host. The capabilities are transferred using the XID channel program.

The primary pieces of new information gained through the MPC XID exchange are:

- MPC to CTC connection
- Polarity of the device (Read or Write)
- I/O buffering information for an active connection

Channel program (X-side or Y-side)

Unlike CTCA, X-Side or Y-Side has meaning only during XID exchange for MPC.

Figure 80 on page 435 shows that how VTAM determines which side is the X-side and which side is the Y-side.

In this example, the operator in subarea 1 is the first to activate the link. Because subarea 2 is not active, the XID exchange does not complete. Later, when the operator in subarea 2 activates the link, the XID exchange is completed.

I/O buffering

During the XID exchange, each host informs the other about its read buffer size.

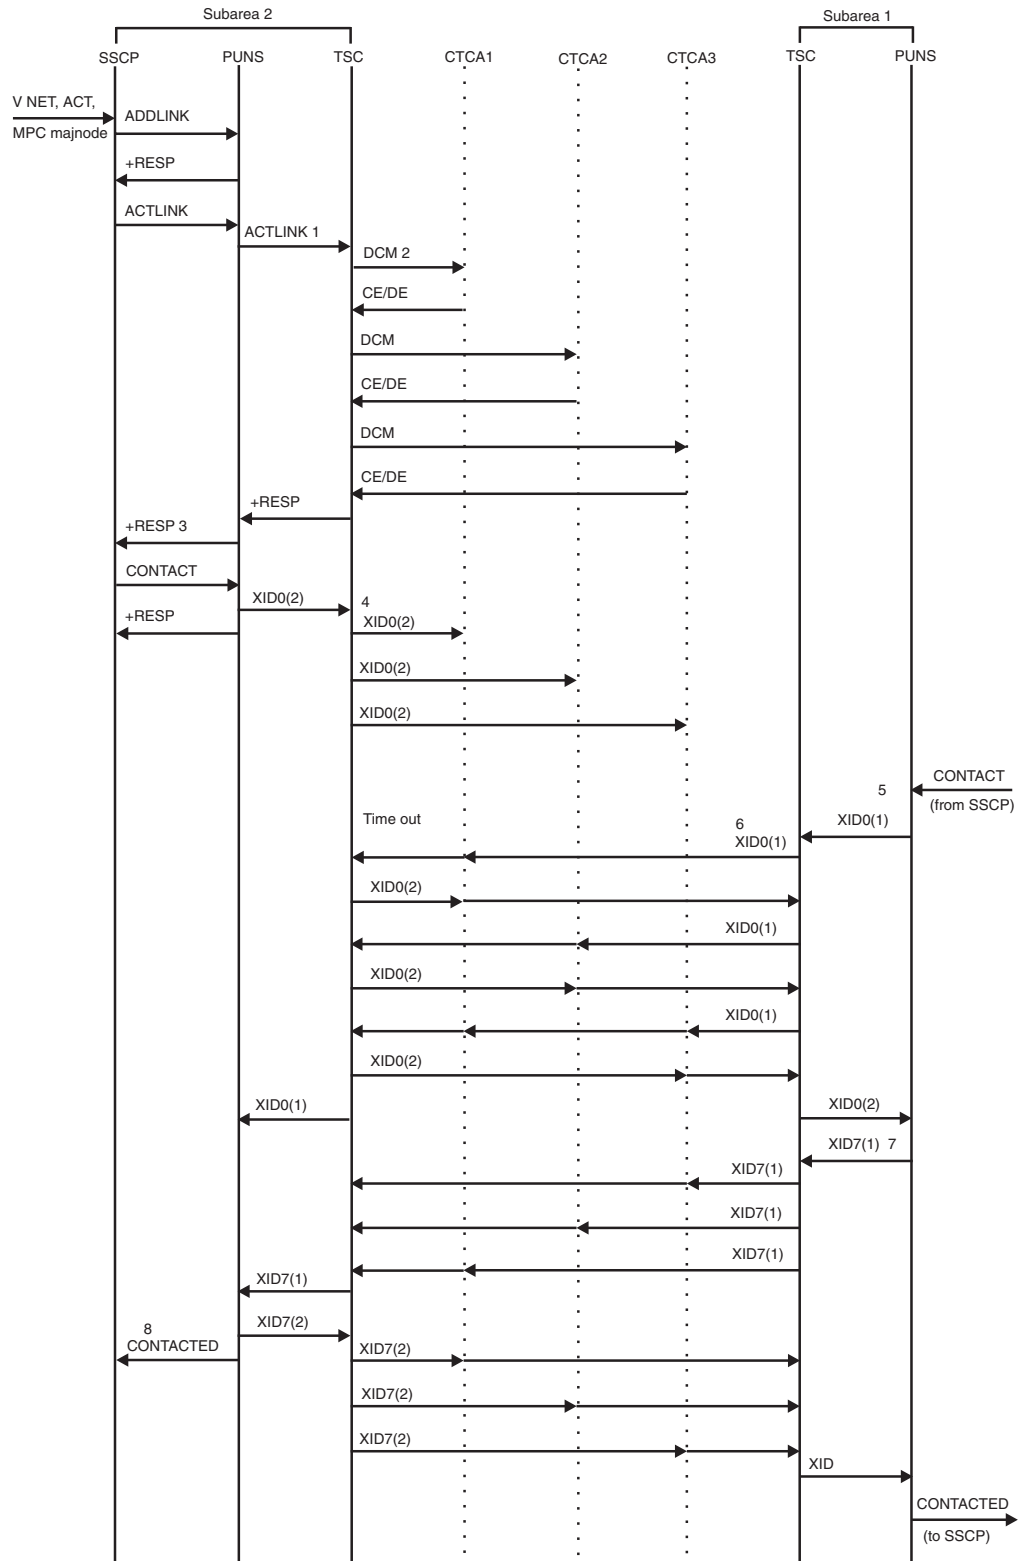


Figure 80. MPC activation flow

CTCA Channel-to-channel adapter
PUNS Physical unit services

SSCP System services control point

TSC Transmission subsystem component

Note: When MPC is used for an APPN host-to-host channel connection, the flows between the SSCP, PUNS, AND TSC components differ somewhat from those in this figure. The XID flows are the same.

1 PUNS queues an ACTLINK RUPE to the MPNCB PUPAB, causing TSC to get control.

2 TSC issues a DCM channel program on every device in the MPC group.

3 A single ACTLINK response is passed to PUNS and SSCP indicating that the MPC path ACTLINK initialization is completed.

4 PUNS XID0 is copied to all MPC subchannels so that they are being written to all MPC subchannels. In this case, subarea 2 XID0s are timed out and the other side is not ready to read.

5 Assume the operator activates the link in subarea 2. Subarea 2 assumes X-side protocols.

6 Subarea 1 initiates writing XID0s at all MPC subchannels and then completes.

7 The subarea with the lower number (subarea 1) responds first to the XID0 exchange by sending an XID7.

8 Subarea 2 checks the XID7 from subarea 1 and responds with its XID7. It also sends a CONTACTED RU to the SSCP.

Channel commands for activating the MPC connection

See “Channel commands for channel-to-channel (CTC) adapters” on page 422 for the commands used for activation and deactivation of the MPC connection.

The MPC XID channel programs are identical to the CTC packed format XID channel programs. See “XID channel program (X-side)” on page 424 and “XID channel program (Y-side)” on page 425.

Channel programs for MPC data transfer

MAXBFRU defines a single read buffer. The write buffer in the other host has the same size allocation. PIUs are blocked for transfer, and are written by a single write, and read by a single read.

VTAM uses three CCWs in its data transfer channel program:

- Write CCW
- Read CCW
- Prepare CCW

MPC unique output control information is in the write buffer and input control information is received in the first 8 bytes of the read buffer.

Normal data transfer for channel programs is a PREP CCW followed by a Write CCW (Write device) or a Prepare CCW followed by a Read CCW (Read device).

For HPDT, sets of these CCWs are chained together in the channel program.

During the XID exchange, each host passes buffer size information for the read portion of its channel programs to other hosts in the network. Each host allocates a write buffer to match the other host read buffer.

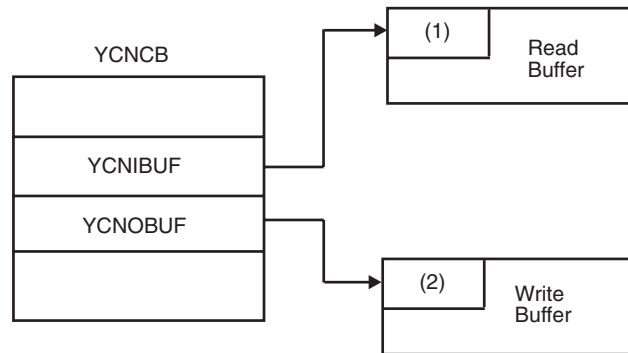


Figure 81. MPC transmit buffers used for normal data transfer

1. The first 8 bytes of the channel buffer is the control information mapped by ISTBKHDR.

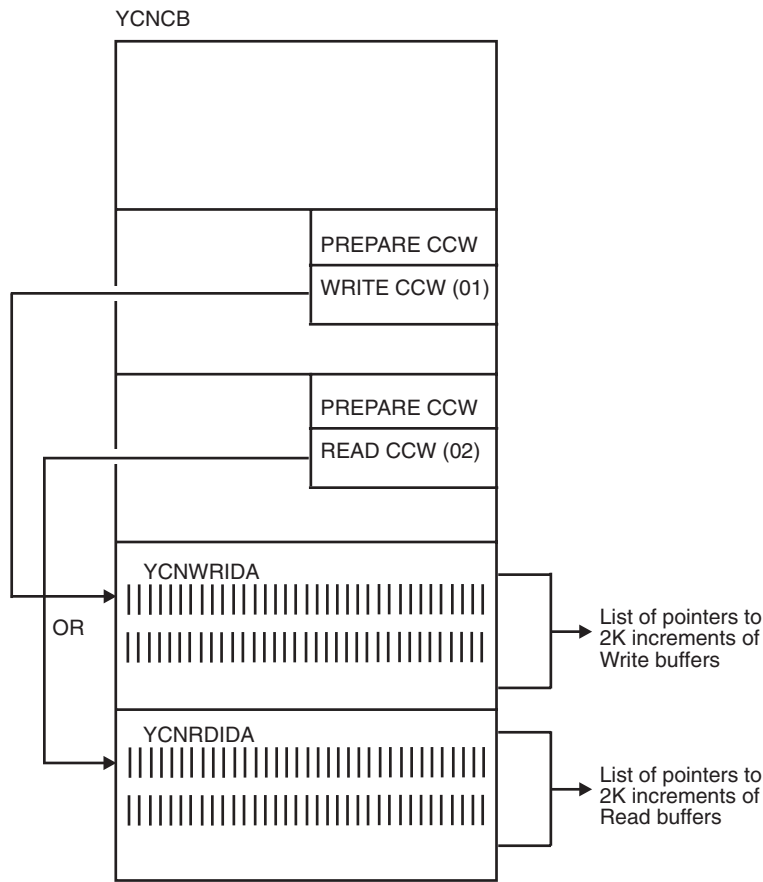


Figure 82. Indirect address word structure of multipath channel programs for normal data flow

Channel program for HPDT MPC data transfer

Seldom Ending Channel Programs (SECP) is for HPDT MPC read and write devices. Because the read and write channel programs can contain up to 17 CCWs, including eight read or write CCWs, the CCWs and the IDAWs are moved out of the YCN CB and into the M2IO and ALPH, respectively. See Figure 83 on page 439 for further details.

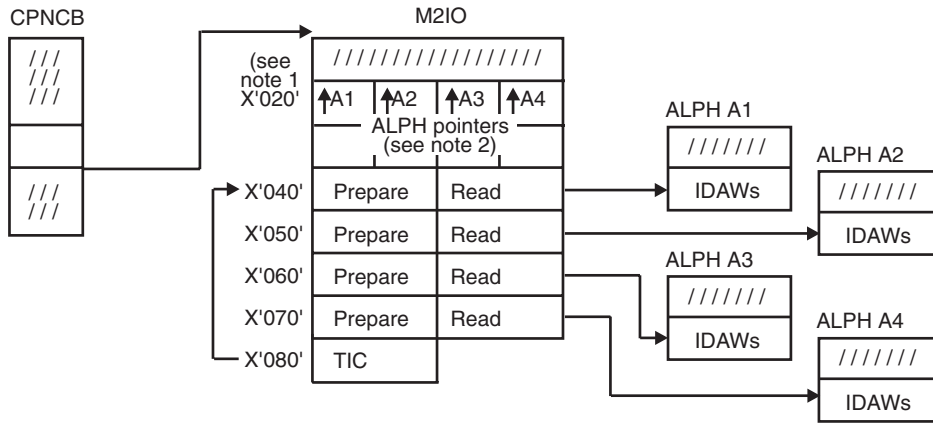


Figure 83. Basic read Seldom Ending Channel Program structure of HPDT multipath channel

Note:

1. M2IO always starts on a page boundary plus X'010' bytes. Offsets shown are from the page boundary.
2. The position of the ALPH pointer corresponds to its CCW set (that is, the ALPH pointer at X'020' contains the address of the ALPH for the prepare/read CCW set at offset X'040'; the ALPH pointer at X'24' contains the address of the ALPH for prepare/read CCW set at offset X'050'; and so on).

01232010	////////	////////	////////	////////	
01232020	012B2800	0129CD98	0129CB20	0128A6E0	← ALPH pointers
01232030	00000000	00000000	00000000	00000000	
01232040	E3600001	01253BF1	0264FFFC	012B29D4	← CCW set 1
01232050	E3600001	01253BF1	0264FFFC	0129CF6C	← CCW set 2
01232060	E3600001	01253BF1	0264FFFC	0129CCF4	← CCW set 3
01232070	E3600001	01253BF1	0224FFFC	0128A8B4	← CCW set 4
01232080	08000000	01232040	////////	////////	
01232090	////////	////////	////////	////////	
012320A0	////////	////////	////////	////////	
012320B0	////////	////////	////////	////////	
012320C0	08000000	01232040			

Figure 84. Example of Seldom Ending Channel Program structure of HPDT multipath channel

SECP read channel program

The read channel program is a logical ring of at least four and at most eight read CCW sets, followed by a TIC CCW back to the physical start of the ring. (A CCW set, when prepare CCWs are required, is composed of both the prepare and read CCWs; otherwise, the set is just the read CCW alone.) The field CPNCB_index_word contains indexes of the logical first and last CCW sets in the ring and is used to maintain control over the channel program and read completion processing. All CCWs are command-chained together except for the last logical CCW in the channel program, which must have command chaining off. PCI is on in all the prepare CCWs, except the CCW in the channel address word when the channel program is started. As reads complete, PCI interrupts occur. The inbound data is processed and routed internally, the read buffers are replenished,

and the CCW set is prepared to be appended to the logical end of the channel program. Preparation is done by ensuring that PCI is on and command chaining is off. The CCW set is appended, or *tacked-in*, to the logical end of the channel program (probably while the channel program is running) by setting the command-chaining bit in the current logical last CCW in the channel program and updating the index of the last logical CCW set in the ring.

The term *missed tack-in* is used to describe the condition where the IO processor fetches the logical last CCW before the CPU sets the command-chaining bit to complete the tack-in. In this case the channel program ends with channel end/device end status. It is the responsibility of the interrupt handler to detect a missed tack-in and to restart the channel program.

The suspend/resume function is not implemented in the read channel program because the objective is to keep as many reads outstanding as required to ensure the channel program does not end.

SECP write channel program

The write channel program is a logical ring of eight write CCW sets followed by a TIC CCW back to the physical start of the ring. However, because write channel programs use suspend/resume, there is a CCW set between the logical end of the ring and the logical start of the ring that cannot be used. Therefore, a maximum of seven writes can be active at one time. The field CPNCB_index_word contains indexes of the logical first and last CCW sets in the ring and is used to maintain control over the channel program and write completion processing. All CCWs are command-chained together and the suspend bit is on in CCW following the last active write. PCI is always off with the exception that it is sometimes set halfway through the channel program when the channel program is large. As writes complete, PCI and suspend interrupts occur. The structures representing the written data are processed, and if another set of data is waiting for transmission, it is tacked-in to the end of the channel program and the channel program is resumed if it is suspended.

As with a read SECP, a missed tack can occur and it is the responsibility of the interrupt handler to resume the channel program.

The suspend/resume function is implemented for write devices because, unlike read devices where reads remain outstanding, it is expected that during normal operation there will be times when there is nothing to write.

Note the following differences between the read and write SECPs:

- The read SECP ALPH pointers in the M2IO are set at link activation time and remain there until the link is deactivated. The write SECP ALPH pointers are nonzero only if the corresponding CCW set is active within the channel program. The write ALPHs can be assigned to any CCW set, while the read ALPHs are always assigned to the same set.
- All 17 CCWs in the M2IO will participate in the IO operations, though not all can be active at the same time.

Channel programs for channel-attached non-SNA 3270 devices

The publications for the non-SNA 3270 devices contain diagnostic procedures. For more information, see the documentation for your display type, or for your control unit. See the z/OS Information Roadmap to determine what document contains information on the 3174 controller.

The LDNCB represents local devices and contains addresses and CCWs needed for channel programs. Figure 85 shows the LDNCB and the location in storage of various CCWs.

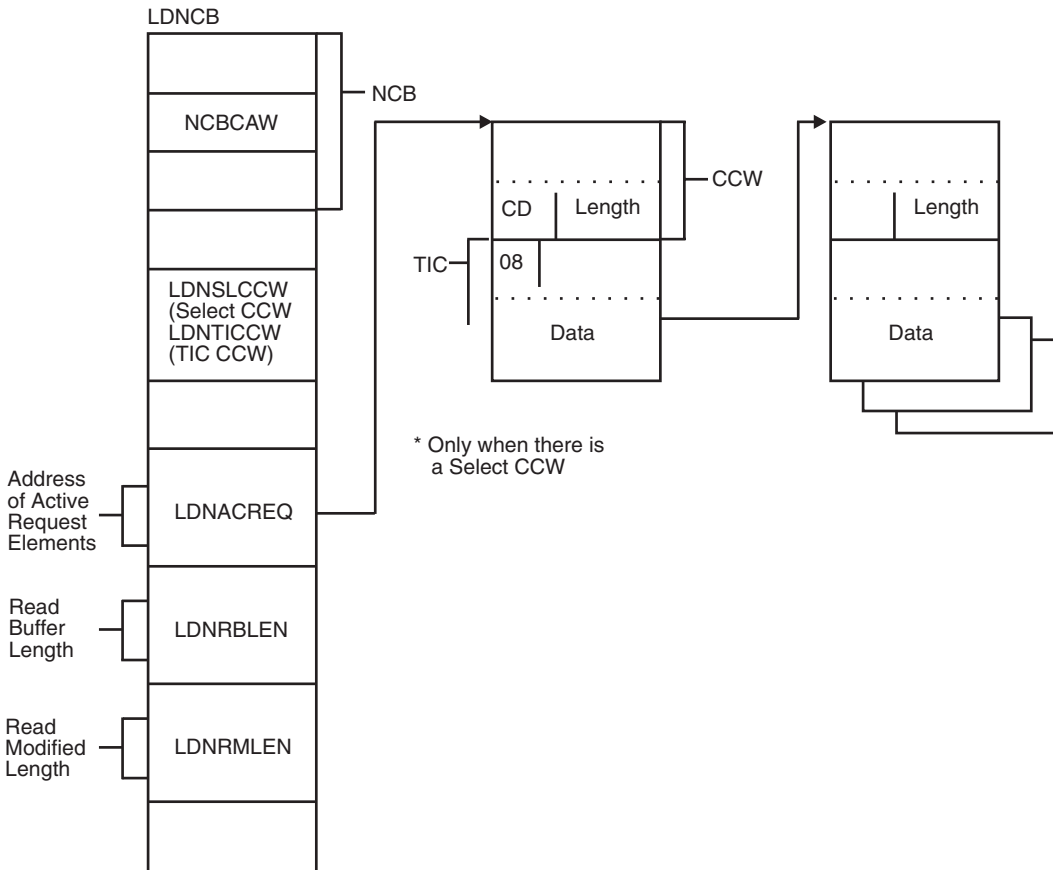


Figure 85. Data areas used by a channel program for channel-attached non-SNA devices

Channel command words

Table 26 contains the channel commands used by VTAM to send data to and receive data from channel-attached non-SNA 3270 terminals.

Table 26. Channel command words for channel-attached non-SNA 3270 devices

Command code (hex)	Command	Description
01	Write	Transfers data from storage in the host processor to the I/O device. Modifies existing buffer data.
02	Read or Read buffer	Transfers the entire buffer contents from the I/O device to storage in the host processor.
05	Erase/Write	Clears the device buffer before starting the write.
06	Read Modified	Transfers data (that was modified since the last read) from the I/O device to storage in the host processor.

Table 26. Channel command words for channel-attached non-SNA 3270 devices (continued)

Command code (hex)	Command	Description
0B	Select	Transfers data from the I/O device buffer to the controller buffer. Separates the device-to-controller unit buffer transfer from the execution command.
0D	erase/write Alternate	Same as Erase/Write, except that it allows for a larger buffer size required for some devices.
1B	Select Read Buffer	Same as Select, except that it is used only for Read buffer channel programs.
4B	Select Write	Same as Select, except that it is used only for Write channel programs.
0F	Erase All Unprotected	Clears all unprotected buffers
11	Write Structured Field	Writes a structured field

Channel programs

Table 27 through Table 32 on page 445 show the CCWs used in various channel programs. The order of execution is the same as the order in which they appear.

Write data channel program

The following table shows how CCWs are used in a Write data channel program.

Table 27. Write data channel program

CCW	Command code (hex)	Flags	Address	Byte count
Select	0B or 4B ⁵	CC,SLI	Zero	1
TIC	08	—	Write CCW	—
Write	01	CD,SLI	Output Area	Length of data
6 (As many as needed)				

Read Modified channel program

The following table shows how CCWs are used in a Read Modified channel program.

Table 28. Read Modified channel program

CCW	Command code (hex)	Flags	Address	Byte count
Select	0B	CC	Zero	1
TIC	08	—	Read Modified CCW	—
Read Modified ⁷	06	CD	Input Area	Length of data
(As many as needed)				
TIC	08	—	Read Modified CCW	—
Read Modified ⁸	06	CD	Input Area	Length of data
Read Modified (Skip)	06	SKIP	Zero	X'7FFF'

Read buffer channel program

The following table shows how CCWs are used in a Read buffer channel program.

Table 29. Read buffer channel program

CCW	Command code (hex)	Flags	Address	Byte count
Select	0B or 1B ⁹	CC,SLI	Zero	1

5. If the UCB indicates that the device will accept the Select Write, code X'4B', it is used until a command reject is received. Then Select, code X'0B', is used.

6. The first Write CCW is pointed to by LDNACREQ.

7. The first Read Modified CCW is pointed to by LDNACREQ.

8. The number of bytes transmitted in the previous Read Modified CCW is stored in the LDNRMLEN field in the LDNCB. For this Read Modified, enough buffers are allocated to hold LDNRMLEN bytes of data. If this is not enough buffers for the device to send all of the pending data, the channel program will end on the Read Skip CCW. The Read Skip CCW reads all of the pending data but does not transmit any of it. The new length is stored in LDNRMLEN. (The length is computed by subtracting the residual byte count in the CSW from X'7FFF'.) Then the channel program is executed again with one more buffer than necessary allocated to read and transmit all of the pending data.

Table 29. Read buffer channel program (continued)

CCW	Command code (hex)	Flags	Address	Byte count
TIC	08	—	Read buffer CCW	—
Read Buffer ¹⁰ (As many as needed)	02	CD	Input Area	Length of data
	TIC	08	Read buffer CCW	—
Read Buffer ¹¹	02	CD	Input Area	Length of data
	02	SKIP	Zero	X'7FFF'
Read Buffer (Skip)				

Erase/Write channel program

The following table shows how CCWs are used in an Erase/Write channel program.

Table 30. Erase/Write channel program

CCW	Command code (hex)	Flags	Address	Byte count
Erase/Write ¹² (As many as needed)	05	CD,SLI	Output Area	Length of data
	TIC	08	Erase/Write CCW	—

9. If the UCB indicates that the device will accept the Select Read, code X'1B', it is used until a command reject is received. Then Select, code X'0B', is used.

10. The first Read CCW is pointed to by LDNACREQ.

11. The number of bytes transmitted in the previous Read buffer CCW is stored in the LDNRBLEN field in the LDNCB. For this Read buffer, enough buffers are allocated to hold LDNRBLEN bytes of data. If this is not enough buffers for the device to send all of the pending data, the channel program will end on the Read Skip CCW. The Read Skip CCW reads all of the pending data but does not transmit any of it. The new length is stored in LDNRBLEN. (The length is computed by subtracting the residual byte count in the CSW from X'7FFF'.) Then the channel program is executed again with enough buffers allocated to read and transmit all of the pending data.

Table 30. Erase/Write channel program (continued)

CCW	Command code (hex)	Flags	Address	Byte count
Erase/Write	05	SLI	Output Area	Length of data

Erase/Write Alternate channel program

The following table shows how CCWs are used in an Erase/Write Alternate channel program.

Table 31. Erase/Write Alternate channel program

CCW	Command code (hex)	Flags	Address	Byte count
Erase/Write ¹³ Alternate (As many as needed)	0D	CD,SLI	Output Area	Length of data
TIC	08	—	Erase/Write Alternate CCW	—
Erase/Write Alternate	0D	SLI	Output Area	Length of data

Erase All Unprotected channel program

The following table shows how CCWs are used in an Erase All Unprotected channel program.

Table 32. Erase All Unprotected channel program

CCW	Command code (hex)	Flags	Address	Byte count
Erase All ¹⁴ Unprotected	0F	SLI	Zero	1

12. LDNACREQ has the address of the first Erase/Write CCW.

13. LDNACREQ contains the address of the first Erase/Write Alternate CCW.

14. LDNACREQ contains the address of the Erase All Unprotected CCW.

Appendix B. Network flows

This appendix describes flows of the VTAM program RUs and AMRUs between network addressable units in single and multiple VTAM networks. Use these flows as guidelines to help analyze and isolate network problems caused by unexpected network events, such as protocol violations. The flow diagrams are divided into the following categories:

- “Generic BIND (GBIND) AMRUs” on page 451
- “Resource activation flows” on page 453
- “Session establishment flows” on page 466
- “Deactivation and session termination flows” on page 492
- “Error detection and recovery and SSCP management services” on page 512

For certain session establishment RUs (ACTCDRM, ACTPU, ACTLU, and BIND), additional RUs can flow if the explicit route (ER) or virtual route (VR) selected for a session is not yet active. Because this flow is essentially the same for all four RUs, these RUs are referred to as generic bind (GBIND) AMRUs. To avoid repetition, the flows for these AMRUs are shown once at the beginning of this appendix.

Table 33 lists all the network flows illustrated in this appendix.

Table 33. Index of network flows

Flow	Page
GBIND AMRU flow	Page
ACTLU: Sending an ACTLU request for a logical unit (LU)	Figure 86 on page 452
ACTPU: Sending an ACTPU request for a communication controller or physical unit (PU)	Figure 87 on page 452
BIND: Sending a BIND request to a secondary logical unit (SLU)	Figure 88 on page 453
Virtual and explicit route: Activating a virtual route (VR) and the associated explicit route (ER)	Figure 89 on page 453
Resource activation flow	Page
Activating a CDRM	
CDRM with COLD response, activating	Figure 104 on page 463
CDRM with ERP response, activating	Figure 103 on page 463
CDRM with a virtual route-based transmission group, activating	Figure 105 on page 463
Activating a cross-network SSCP-SSCP session	
Back-to-back gateway NCPs request sessions	Figure 106 on page 464
Gateway VTAM requests session	Figure 107 on page 465
Non-gateway VTAM requests session	Figure 108 on page 465
Activating an NCP major node	

Table 33. Index of network flows (continued)

Flow	Page
Channel-attached communication controller, activating	Figure 90 on page 455
Link-attached communication controller, activating	Figure 91 on page 456
Activating resources controlled by a host or NCP major node	
Link: Activating a link	Figure 92 on page 456
Link station: Activating a cross-subarea link station	Figure 94 on page 457
Logical unit (LU): Activating a logical unit	Figure 101 on page 462
Application program: Activating an application program and processing an OPEN ACB request	Figure 102 on page 462
Physical unit (PU): Activating a physical unit type 2.0	Figure 96 on page 459
Physical unit (PU): Activating a physical unit type 2.0 with load required	Figure 97 on page 459
Physical unit (PU): Moving a dynamically added physical unit	Figure 99 on page 461
Physical unit (PU): Moving a SYSGENed physical unit	Figure 98 on page 460
SSCP takeover of peripheral node logical units (LUs)	Figure 100 on page 461
Switched connection, establishing	Figure 95 on page 458
Switched link with takeover, activating	Figure 93 on page 457
Session establishment flow	Page
Failed session establishment	
Failure (CDINIT rejection) of session initiation by a secondary logical unit (SLU) for single gateway VTAM and single gateway NCP	Figure 134 on page 491
Failure (CINIT rejection) of setup procedure initiated by a secondary logical unit (SLU) for single gateway VTAM and single gateway NCP	Figure 136 on page 492
Failure (SETCV failure) of session initiation by a secondary logical unit (SLU) for single gateway VTAM and single gateway NCP	Figure 135 on page 491
Gateway VTAM	
Default partitioning of gateway VTAM responsibility spanning three networks	Figure 127 on page 484
Multiple gateway VTAMs and back-to-back gateway NCPs	Figure 110 on page 469
PLU availability for autologon, notification of	Figure 133 on page 490
Requests initiated by primary logical units (PLUs)	
Dependent PLU initiating cross-domain session with independent SLU	Figure 116 on page 474
Independent PLU initiating cross-domain session with independent SLU	Figure 115 on page 474
Independent PLU requesting session with independent SLU through a single gateway VTAM and single gateway NCP	Figure 118 on page 476
OPNDST ACQUIRE	Figure 111 on page 470
PLU initiating request for single gateway VTAM and single gateway NCP	Figure 117 on page 475
PLU initiating request setup queued for single gateway VTAM and single gateway NCP	Figure 119 on page 477
SIMLOGON	Figure 112 on page 471
SIMLOGON(RELREQ)	Figure 113 on page 472
SIMLOGON(RELREQ): Session is pending active or already in progress	Figure 114 on page 473

Table 33. Index of network flows (continued)

Flow	Page
Requests initiated by secondary logical units (SLUs)	
Dependent SLU initiating cross-domain session with application logical unit (LU)	Figure 124 on page 481
INIT SELF	Figure 122 on page 480
LOGON	Figure 120 on page 478
Predesignated control of gateway NCP by middle host	Figure 127 on page 484
REQSESS	Figure 121 on page 479
Sending an unformatted request to the SSCP	Figure 123 on page 480
Single gateway connecting three or more networks	Figure 126 on page 483
Single gateway VTAM and single gateway NCP	Figure 125 on page 482
Requests initiated by third parties	
CLSDST PASS	Figure 129 on page 486
CLSDST PASS with NOTIFY	Figure 130 on page 487
Request spanning three networks	Figure 131 on page 488
VARY LOGON or LOGAPPL processing	Figure 132 on page 489
Deactivation or session termination flow	
Page	
CLOSE ACB processing	Figure 152 on page 504
Deactivating an application program	Figure 153 on page 505
Deactivating a CDRM	
Forced	Figure 156 on page 507
Forced, without affecting active sessions	Figure 158 on page 508
Forced or immediate, VTAM releases before V3R4.1	Figure 159 on page 509
Immediate	Figure 155 on page 507
Immediate, without affecting active sessions	Figure 157 on page 508
Normal	Figure 154 on page 506
Deactivating a logical unit (LU), single network	
Forced	Figure 138 on page 494
Immediate	Figure 137 on page 494
VARY NET,TERM Cleanup	Figure 150 on page 503
VARY NET,TERM Unconditional	Figure 149 on page 502
With Giveback	Figure 139 on page 495
Deactivating a logical unit (LU), multiple networks	
Independent PLU sends BFCLEANUP for independent SLU	Figure 140 on page 495
Independent PLU sends UNBIND for independent SLU	Figure 141 on page 496
PLU sends UNBIND for multiple gateway VTAMs and single gateway NCP	Figure 142 on page 496
PLU sends UNBIND for single gateway VTAM and single gateway NCP	Figure 143 on page 497
SLU requests TERMINATE SELF (CLEANUP) for single gateway VTAM and single gateway NCP	Figure 145 on page 499

Table 33. Index of network flows (continued)

Flow	Page
SLU requests TERMINATE SELF for multiple gateway VTAMs and back-to-back gateway NCPs	Figure 144 on page 498
SLU requests TERMINATE SELF for single gateway VTAM and single gateway NCP	Figure 146 on page 500
Type 2.1 nodes, active termination	Figure 147 on page 501
Deactivating a physical unit (PU) acting as an adjacent link station for independent logical unit (LU) sessions	Figure 148 on page 501
Queued session, terminating	Figure 151 on page 503
Route failure	
Route failure in intermediate network causes termination of LU-LU sessions	Figure 161 on page 511
Route failure in intermediate network causes termination of SSCP-SSCP sessions	Figure 162 on page 511
SSCP-SSCP session termination causes LU-LU sessions to be broken	Figure 160 on page 510
Error and SSCP management services flow	Page
Error recovery processing (ERP)	
Hard INOP	Figure 164 on page 513
Soft INOP	Figure 163 on page 513
LPDA-2 processing	
Unsolicited LPDA-2 test on permanent link error with two link segments	Figure 168 on page 516
Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 physical unit (PU) with one link segment	Figure 166 on page 514
Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 physical unit (PU) with two link segments	Figure 167 on page 515
SSCP management services processing	
FORWARD and DELIVER Routing	Figure 165 on page 514
XRF processing	
Secondary logical unit (LU) initiate with USERVAR (LOGON)	Figure 171 on page 519
Third-party initiate (CLSDST PASS)	Figure 172 on page 520
XRF primary and backup sessions, establishment of	Figure 169 on page 517
XRF session switch (takeover)	Figure 170 on page 518

Many abbreviations are shown at the top of the network flows. The following list gives the meaning of some of those abbreviations:

- APPL** Application
- BF** Boundary function
- BFSS** Boundary function session services
- BNN** Boundary network node
- CS** Configuration services
- EU** End user

LU	Logical unit
NCP	Network Control Program
NOS	Network operator services
PLU	Primary logical unit
PN	Peripheral node
PU	Physical unit
PUNS	Physical unit services
SLU	Secondary logical unit
SS	Session services
SSCP	System services control point
TSC	Transmission subsystem component
XRF	Extended recovery facility

Generic BIND (GBIND) AMRUs

Access method RUs (AMRUs) are internal requests that might appear in the PIU trace and are a function of physical unit services (PUNS), configuration services, or session services.

Figure 86 on page 452 through Figure 89 on page 453 show the flow of these requests and responses between the SSCP and logical and physical units when a virtual route (VR) or explicit route (ER) selected for a session is not yet active.

Index of generic BIND (GBIND) AMRU flows

Table 34 lists the GBIND AMRU flows illustrated here.

Table 34. Index of generic BIND (GBIND) AMRU flows

Flow	Page
ACTLU: Sending an ACTLU request for a logical unit (LU)	Figure 86 on page 452
ACTPU: Sending an ACTPU request for a communication controller or physical unit (PU)	Figure 87 on page 452
BIND: Sending a BIND request to a secondary logical unit (SLU)	Figure 88 on page 453
Virtual and explicit route: Activating a virtual route (VR) and the associated explicit route (ER)	Figure 89 on page 453

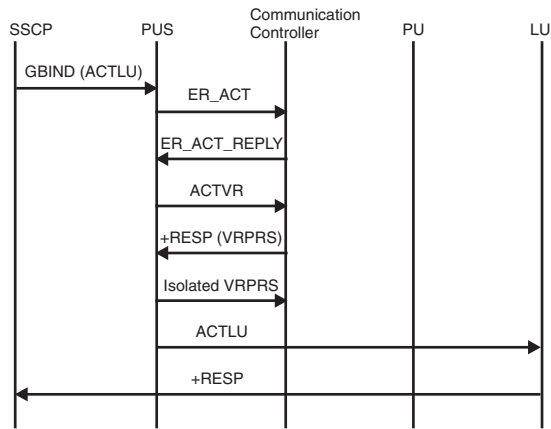


Figure 86. Sending an ACTLU request for a logical unit (LU)

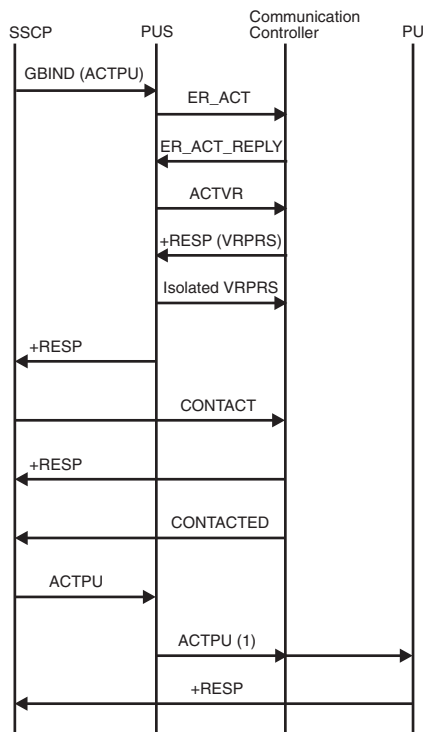


Figure 87. Sending an ACTPU request for a communication controller or physical unit (PU)

1. The ACTPU can flow either to the communication controller or to a physical unit.

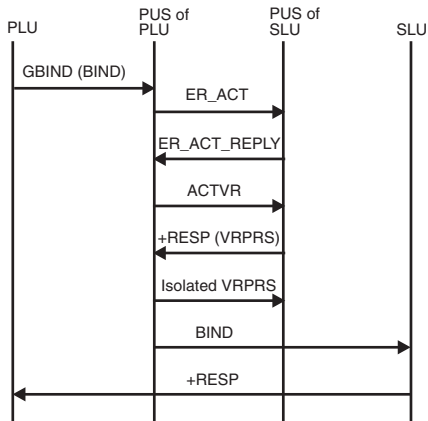


Figure 88. Sending a BIND request to a secondary logical unit (SLU)

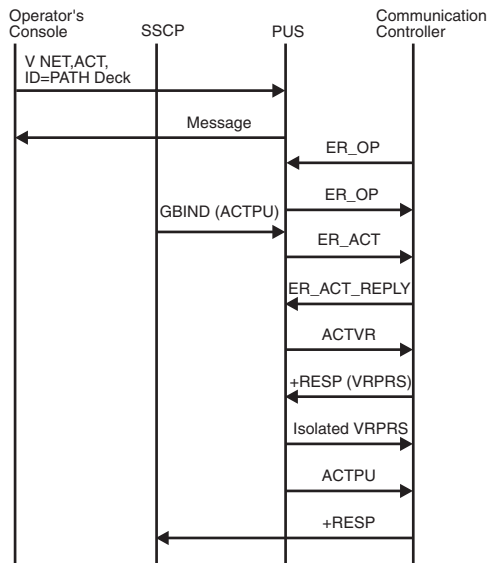


Figure 89. Activating a virtual route (VR) and the associated explicit route (ER)

Resource activation flows

Figure 90 on page 455 through Figure 108 on page 465 show the flow of requests and responses between the SSCP and logical and physical units to activate resources.

For channel activation flows, see Appendix A, "Channel programs," on page 413.

Index of resource activation flows

Table 35 lists the resource activation flows that are illustrated here.

Table 35. Index of resource activation flows

Flow	Page
Activating a CDRM	

Table 35. Index of resource activation flows (continued)

Flow	Page
CDRM with COLD response, activating	Figure 104 on page 463
CDRM with ERP response, activating	Figure 103 on page 463
CDRM with a virtual route-based transmission group, activating	Figure 105 on page 463
Activating a cross-network SSCP-SSCP session	
Back-to-back gateway NCPs request sessions	Figure 106 on page 464
Gateway VTAM requests session	Figure 107 on page 465
Non-gateway VTAM requests session	Figure 108 on page 465
Activating an NCP major node	
Channel-attached communication controller, activating	Figure 90 on page 455
Link-attached communication controller, activating	Figure 91 on page 456
Activating resources controlled by a host or NCP major node	
Link: Activating a link	Figure 92 on page 456
Link station: Activating a cross-subarea link station	Figure 94 on page 457
Logical unit (LU): Activating a logical unit	Figure 101 on page 462
Application program: Activating an application program and processing an OPEN ACB request	Figure 102 on page 462
Physical unit (PU): Activating a physical unit type 2.0	Figure 96 on page 459
Physical unit (PU): Activating a physical unit type 2.0 with load required	Figure 97 on page 459
Physical unit (PU): Moving a dynamically added physical unit	Figure 99 on page 461
Physical unit (PU): Moving a SYSGENed physical unit	Figure 98 on page 460
SSCP takeover of peripheral node logical units (LUs)	Figure 100 on page 461
Switched connection, establishing	Figure 95 on page 458
Switched link with takeover, activating	Figure 93 on page 457

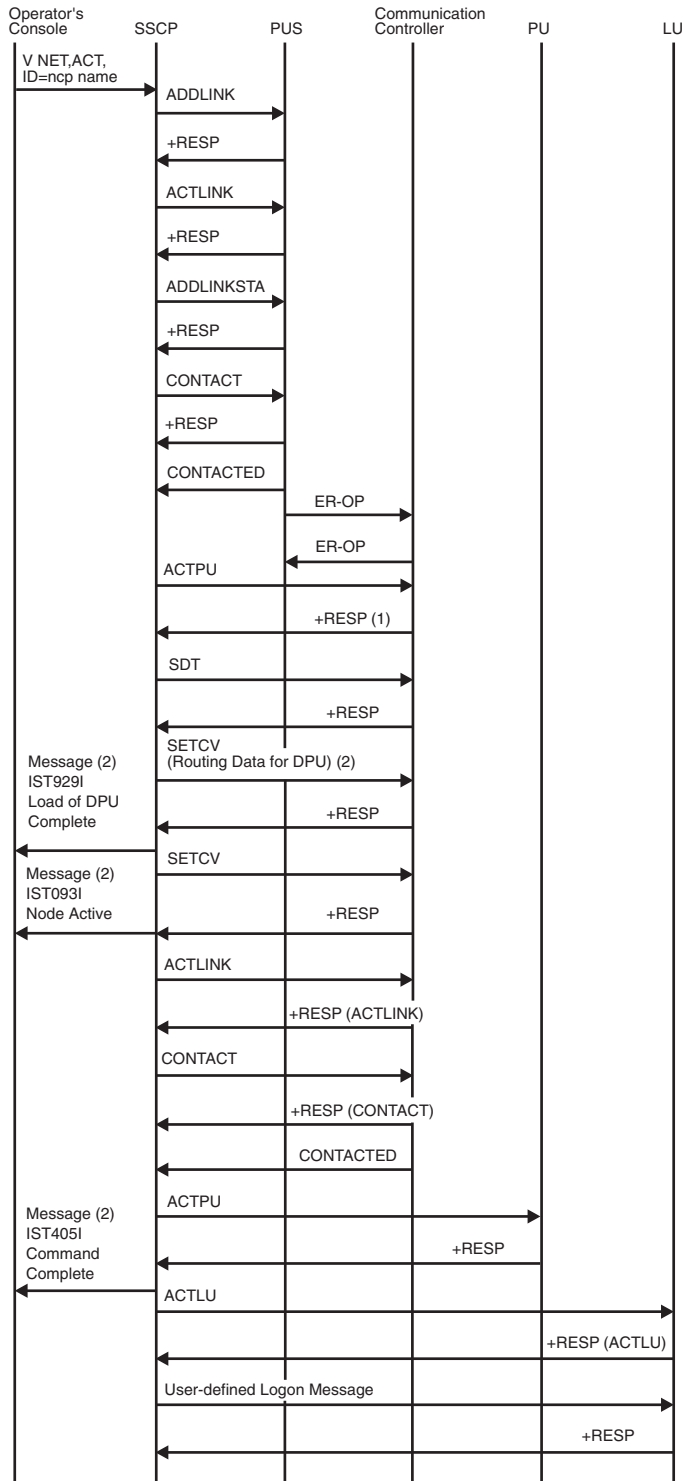


Figure 90. Activating a channel-attached communication controller

1. Includes NCP dynamic path definition capability indicator.
2. Flows only for dynamic path definition. SETCV and IST929I flow for each dynamic path definition member specified.

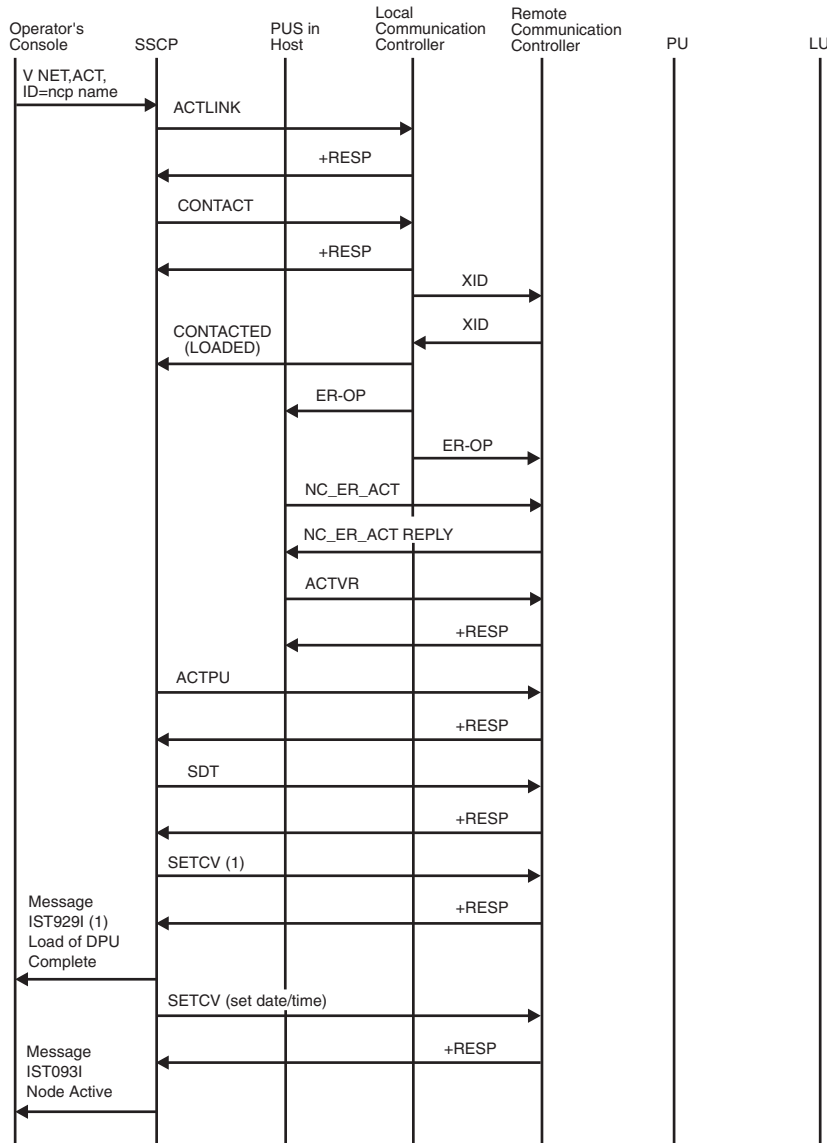


Figure 91. Activating a link-attached communication controller

1. Flows only for dynamic path definition. SETCV and IST929I flow for each dynamic path definition member specified.




Figure 92. Activating a link (ACTLINK)

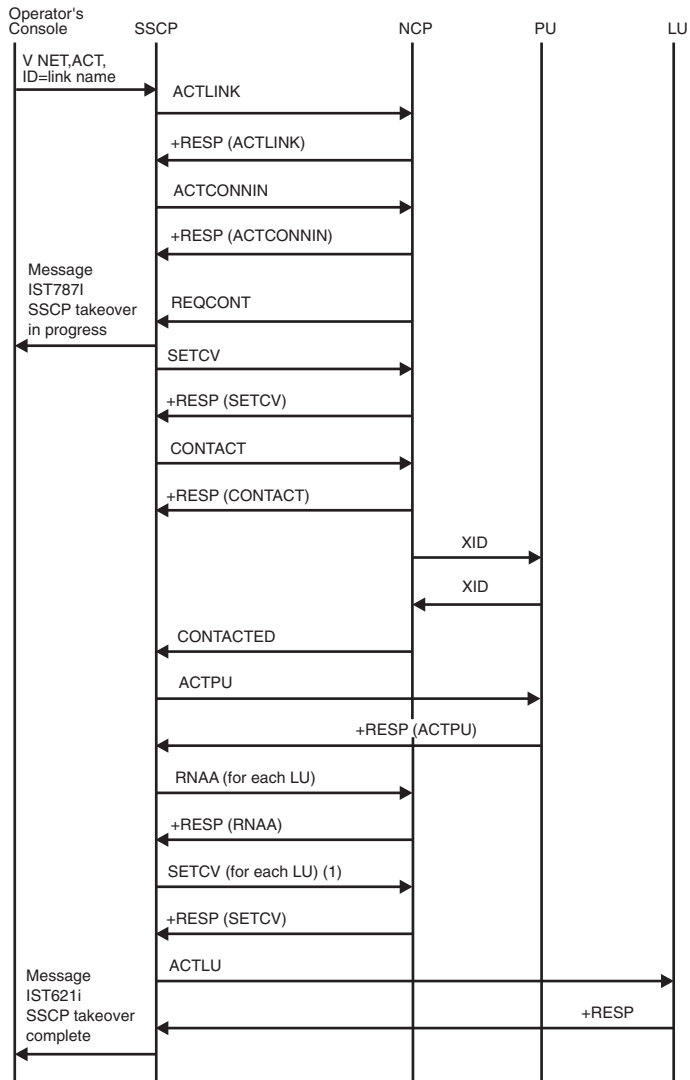


Figure 93. Activating a switched link with takeover

1. SETCV does not flow for NCPs that support peripheral nodes.

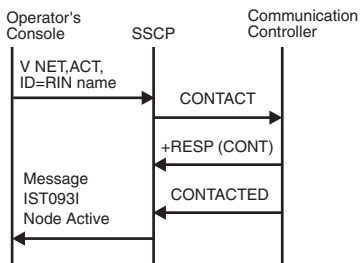


Figure 94. Activating a cross-subarea link station

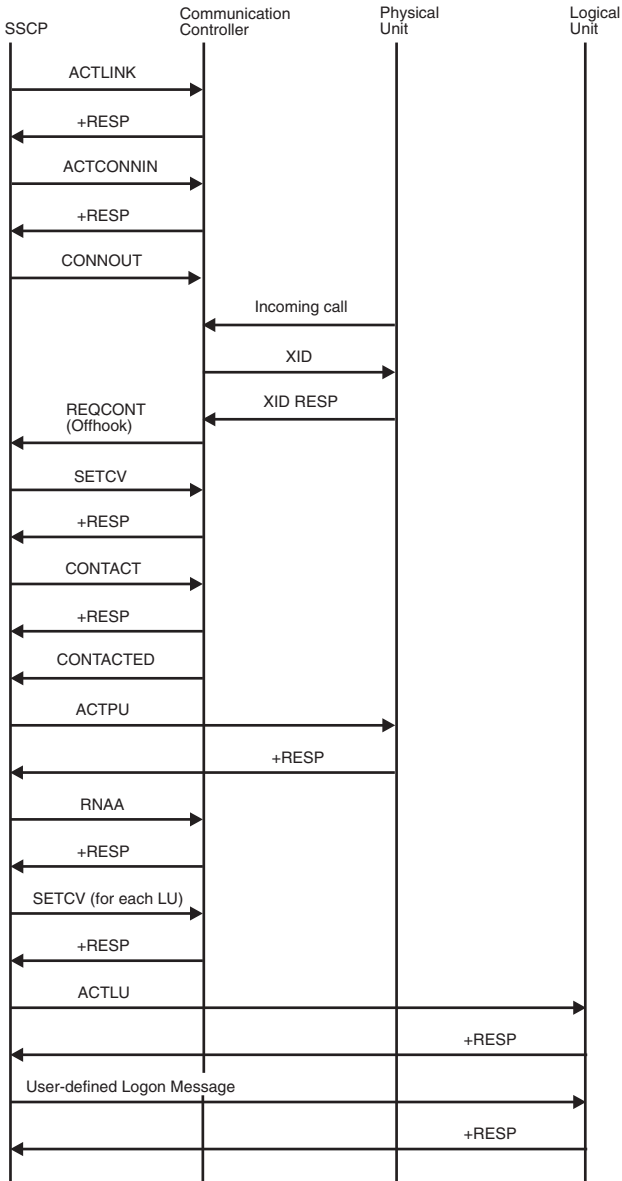


Figure 95. Establishing a switched connection

To establish a switched connection, the SSCP sends an Activate Link request to indicate that the link is active. An Activate Connect In request is sent to enable the communication controller to answer incoming calls. (Instead of Activate Connect In, Dial could be sent to initiate an outbound call.) When a call comes in, the communication controller sends an exchange identification (XID) and the physical unit responds with its ID (station address). The communication controller sends a Request Contact (Offhook) request to the SSCP. The SSCP sends a Set Control Vector request containing address and pacing information to the communication controller. The standard activation sequence then occurs.

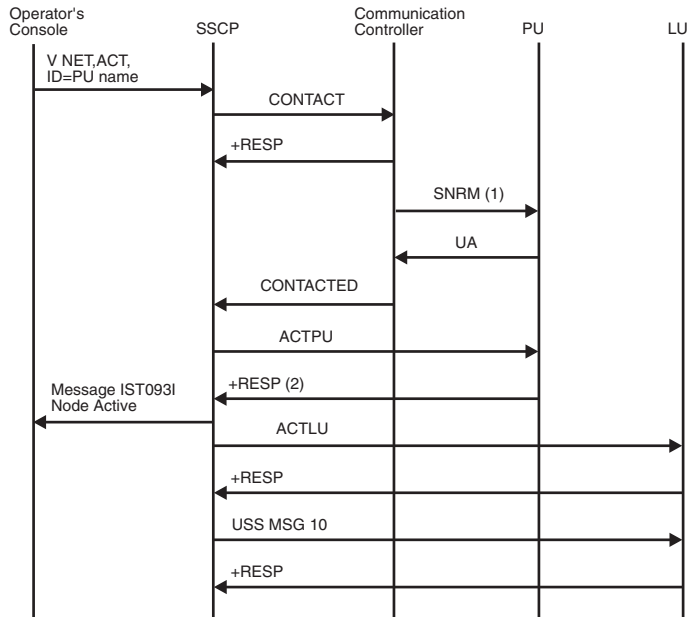


Figure 96. Activating a physical unit type 2.0

1. An XID, instead of an SNRM, will flow to a switched line.
2. Additional RUs flow if the physical unit must be loaded. These RUs are shown in Figure 97.

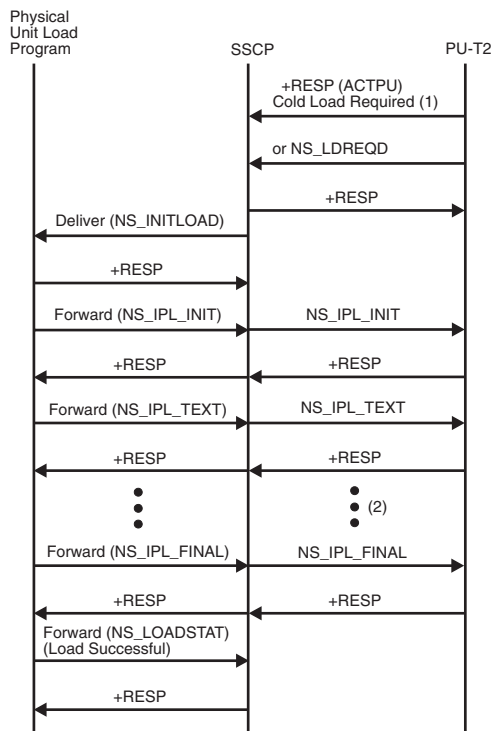


Figure 97. Activating a physical unit type 2.0 with load required

1. This figure shows only the RUs that flow when a type 2 physical unit requires loading. For the RUs that flow before and after those in this figure, see Figure 96 on page 459.
2. NS_IPL_TEXT and the response might repeat.

For type 2.0 physical units that require loading before they can be activated, the request for load is indicated in the ACTPU response. (During the activation, the physical unit might request loading with an NS_LDREQD RU.)

The SSCP formats the load request into a network services (NS) RU to initiate the load. The management services subcomponent of the SSCP then sends the embedded request to the physical unit load program of the Downstream Load Utility.

If the physical unit load program *is not* available, it sends a negative response to the SSCP's Deliver RU. The SSCP then sends an NS_IPL_ABORT RU to the physical unit for deactivation processing. (If the load was requested with an NS_LDREQD RU, the physical unit is not deactivated; in fact, it might try the load request again.)

If the physical unit load program *is* available, as in Figure 97 on page 459, it sends a positive response to the SSCP's Deliver RU. When the load program is complete, it sends a Forward RU, containing an NS_LOADSTAT RU, to relay the status of the load operation.

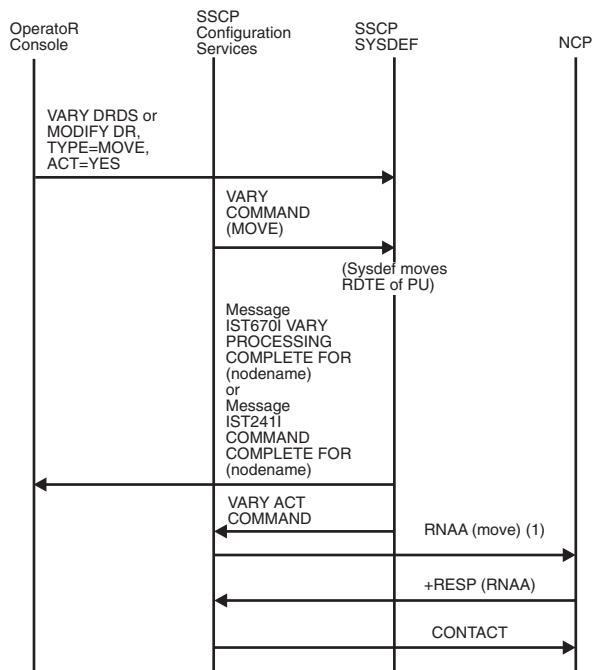


Figure 98. Moving a SYSGENed physical unit

1. RNAA flow is as normal. RNAA does not flow for MODIFY DR, TYPE=MOVE, ACT=NO.

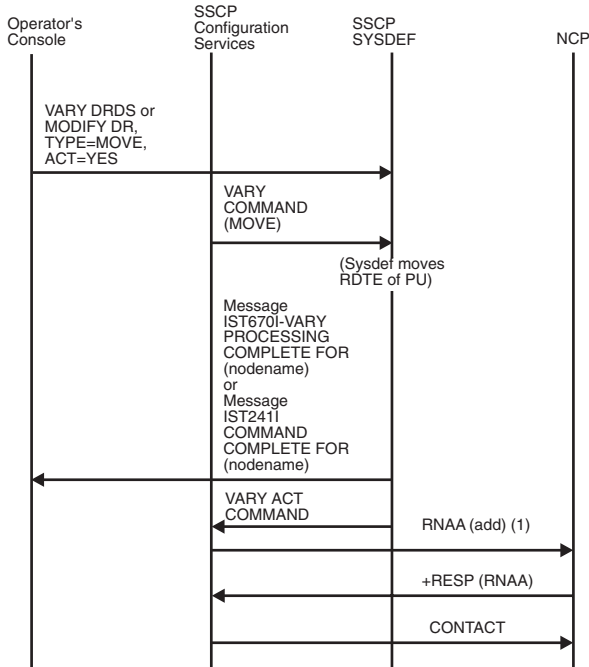


Figure 99. Moving a dynamically added physical unit

1. RNAA flow is as normal. RNAA does not flow for MODIFY DR, TYPE=MOVE, ACT=NO.

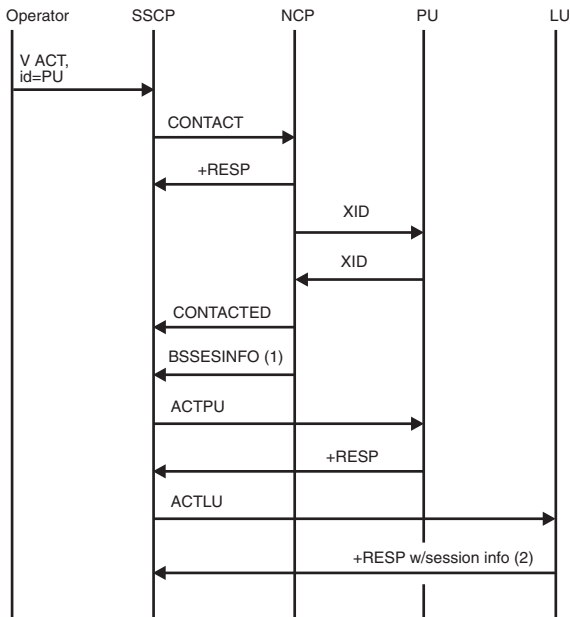


Figure 100. SSCP takeover of peripheral node logical units

Note: The following conditions are assumed for this example:

1. The physical units being taken over are defined with ANS=CONTINUE, specifying that any LU-LU sessions that are active at SSCP-failure time will continue.
2. There are some LU-LU sessions active at failure time under the physical unit being taken over.
3. Independent logical unit only (possible multiple RUs).
4. Dependent logical unit only.

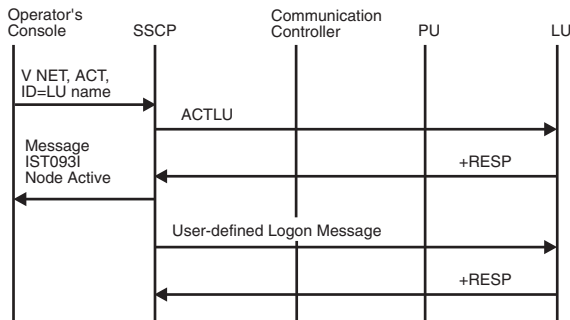


Figure 101. Activating a logical unit

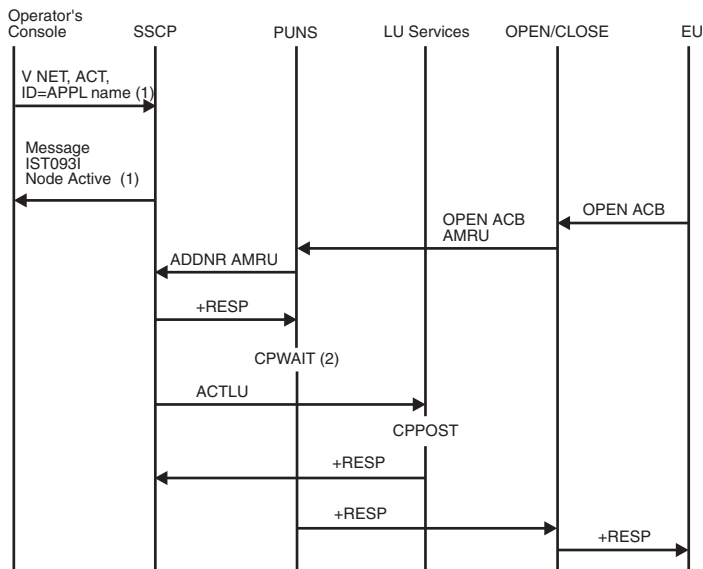


Figure 102. Activating an application program and processing an OPEN ACB request

1. These do not flow for OPEN ACB processing.
2. PUNS cannot send a response to the OPEN ACB request until LUS receives an ACTLU request for the application program. Therefore, PUNS issues CPWAIT and waits for LUS to post it. After LUS has received the ACTLU, it posts PUNS, which then sends a response to the OPEN ACB request.

For the close ACB flow, see Figure 152 on page 504.

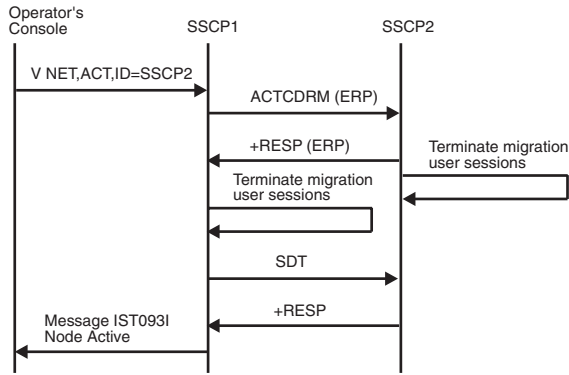


Figure 103. Activating CDRM with ERP response

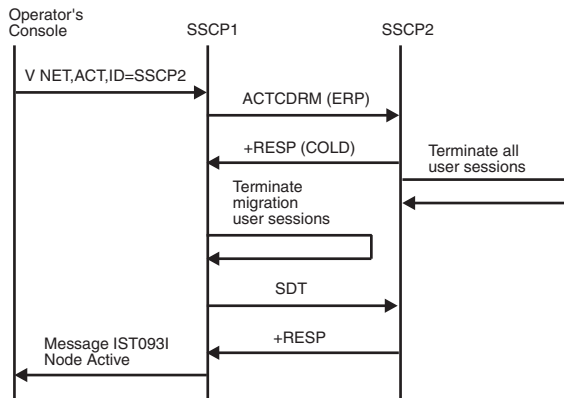


Figure 104. Activating CDRM with COLD response

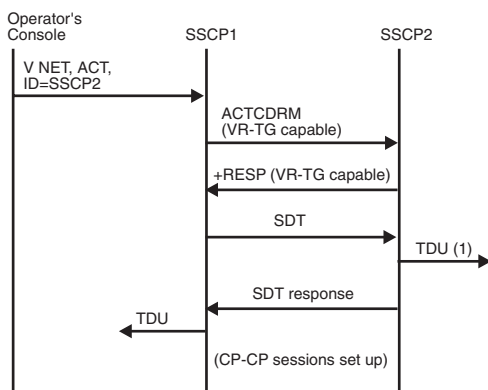


Figure 105. Activating a CDRM with a virtual-route-based transmission group

1. If the transmission group is an intermediate routing transmission group (NN-NN), the topology database update (TDU) will be built and broadcast. If the host is a migration data host, the topology database update (TDU) will be built and sent to its server.

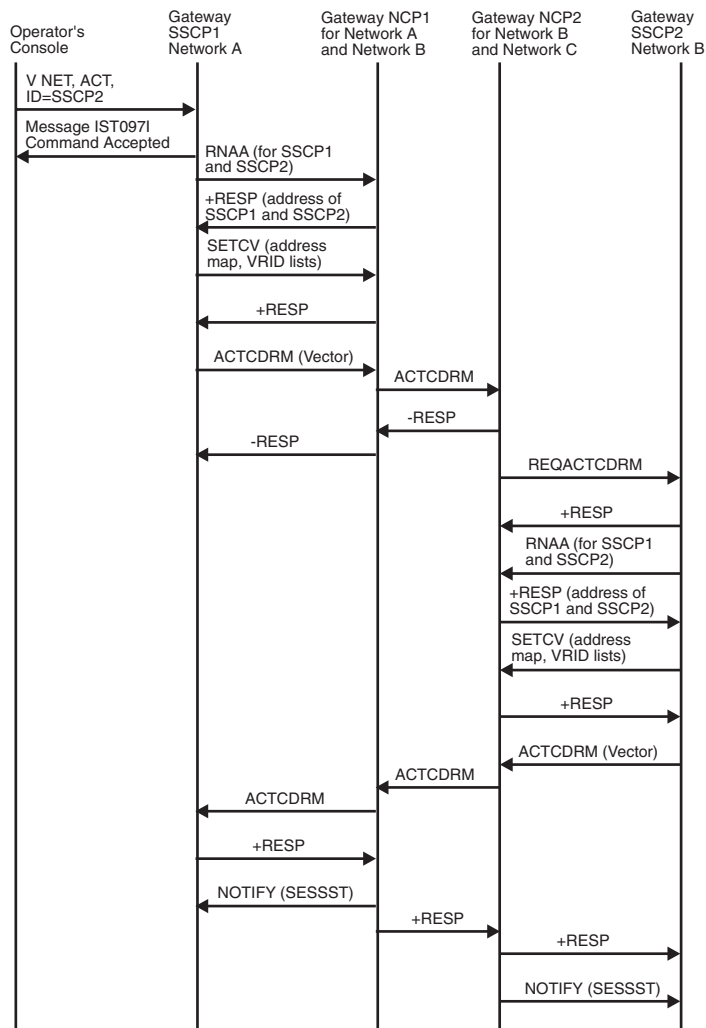


Figure 106. Back-to-back gateway NCP request sessions

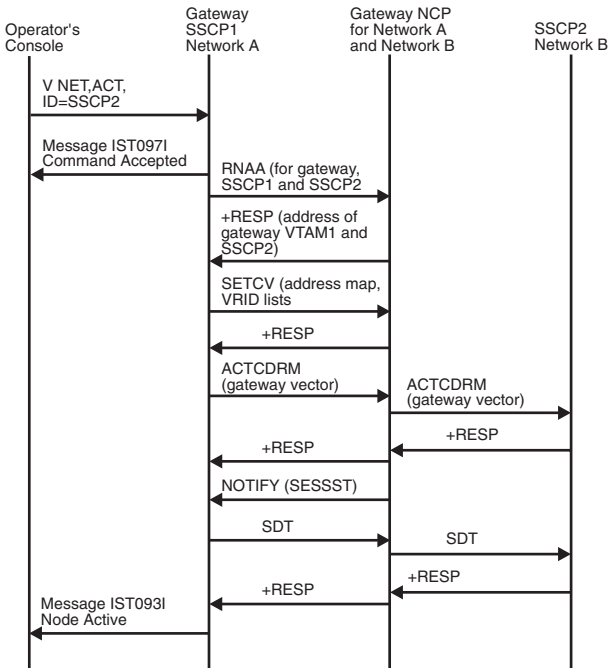


Figure 107. Gateway VTAM requests session

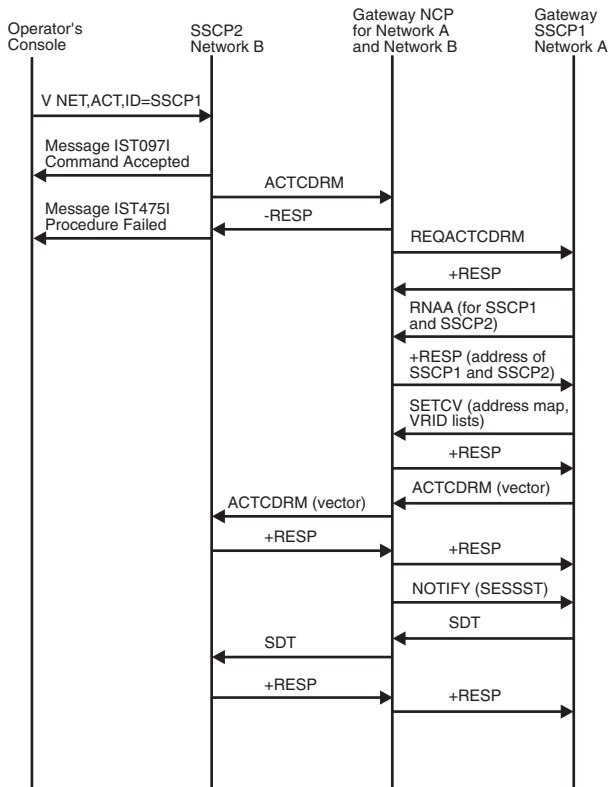


Figure 108. Non-gateway VTAM requests session

Session establishment flows

Figure 109 on page 468 through Figure 136 on page 492 show the flow of requests and responses to establish single network and cross-network LU-LU sessions.

Index of session establishment flows

Table 36 lists the session establishment flows that are illustrated here.

Table 36. Index of session establishment flows

Flow	Page
Failed session establishment	
Failure (CDINIT rejection) of session initiation by a secondary logical unit (SLU) for single gateway VTAM and single gateway NCP	Figure 134 on page 491
Failure (CINIT rejection) of setup procedure initiated by a secondary logical unit (SLU) for single gateway VTAM and single gateway NCP	Figure 136 on page 492
Failure (SETCV failure) of session initiation by a secondary logical unit (SLU) for single gateway VTAM and single gateway NCP	Figure 135 on page 491
Gateway VTAM	
Default partitioning of gateway VTAM responsibility spanning three networks	Figure 127 on page 484
Multiple gateway VTAMs and back-to-back gateway NCPs	Figure 110 on page 469
PLU availability for autologon, notification of	
Requests initiated by primary logical units (PLUs)	
Dependent PLU initiating cross-domain session with independent SLU	Figure 116 on page 474
Independent PLU initiating cross-domain session with independent SLU	Figure 115 on page 474
Independent PLU requesting session with independent SLU through a single gateway VTAM and single gateway NCP	Figure 118 on page 476
OPNDST ACQUIRE	Figure 111 on page 470
PLU initiating request for single gateway VTAM and single gateway NCP	Figure 117 on page 475
PLU initiating request setup queued for single gateway VTAM and single gateway NCP	Figure 119 on page 477
SIMLOGON	Figure 112 on page 471
SIMLOGON(RELREQ)	Figure 113 on page 472
SIMLOGON(RELREQ): Session is pending active or already in progress	Figure 114 on page 473
Requests initiated by secondary logical units (SLUs)	
Dependent SLU initiating cross-domain session with application logical unit (LU)	Figure 124 on page 481
INIT SELF	Figure 122 on page 480
LOGON	Figure 120 on page 478
Pre-designated control of gateway NCP by middle host	Figure 127 on page 484
REQSESS	Figure 121 on page 479
Sending an unformatted request to the SSCP	Figure 123 on page 480
Single gateway connecting three or more networks	Figure 126 on page 483
Single gateway VTAM and single gateway NCP	Figure 125 on page 482

Table 36. Index of session establishment flows (continued)

Flow	Page
Requests initiated by third parties	
CLSDST PASS	Figure 129 on page 486
CLSDST PASS with NOTIFY	Figure 130 on page 487
Request spanning three networks	Figure 131 on page 488
VARY LOGON or LOGAPPL processing	Figure 132 on page 489

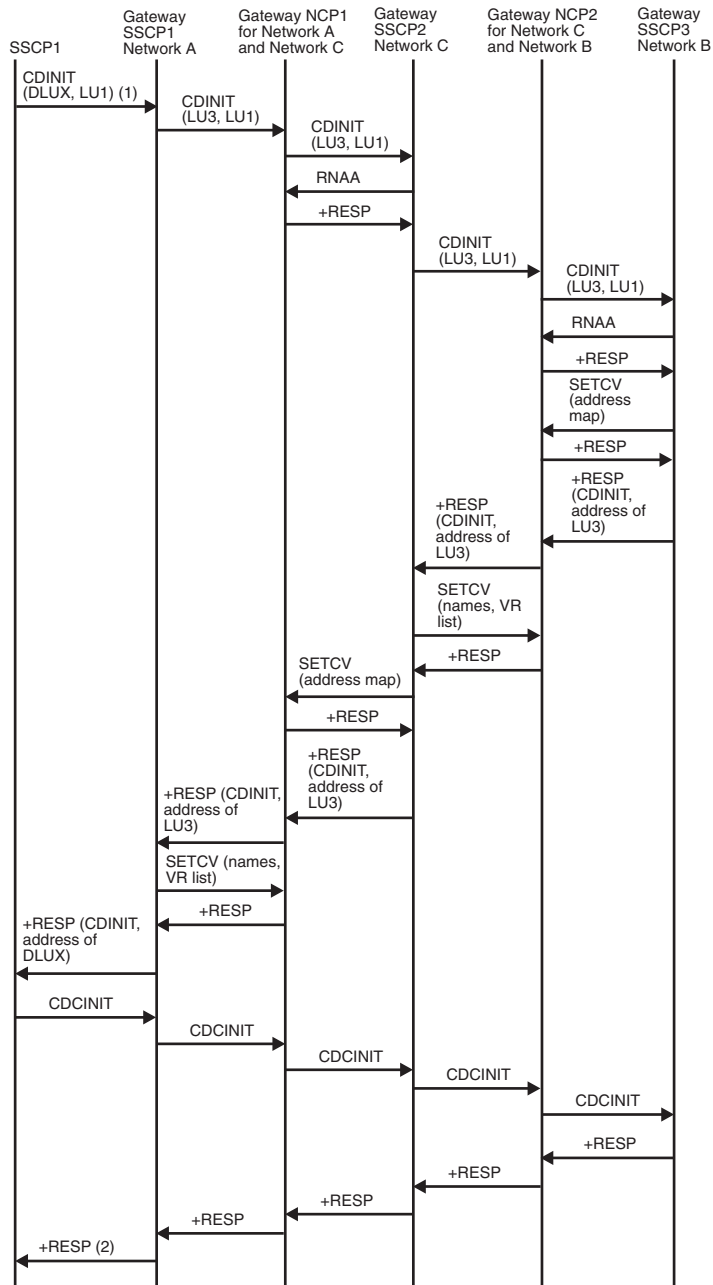


Figure 109. Default partitioning of gateway VTAM responsibility spanning three networks

1. Only the SSCP-SSCP session communication is shown. Assume LU1 (a logical unit owned by SSCP1 in Network A) requests a session with DLUX (an alias for LU3 in Network B; LU3 is a logical unit owned by gateway VTAM3). LU1 is the SLU in the request session.
2. Session setup proceeds as shown in the basic flows. BIND flows from LU3 to gateway NCP2, to gateway NCP1, and to LU1.

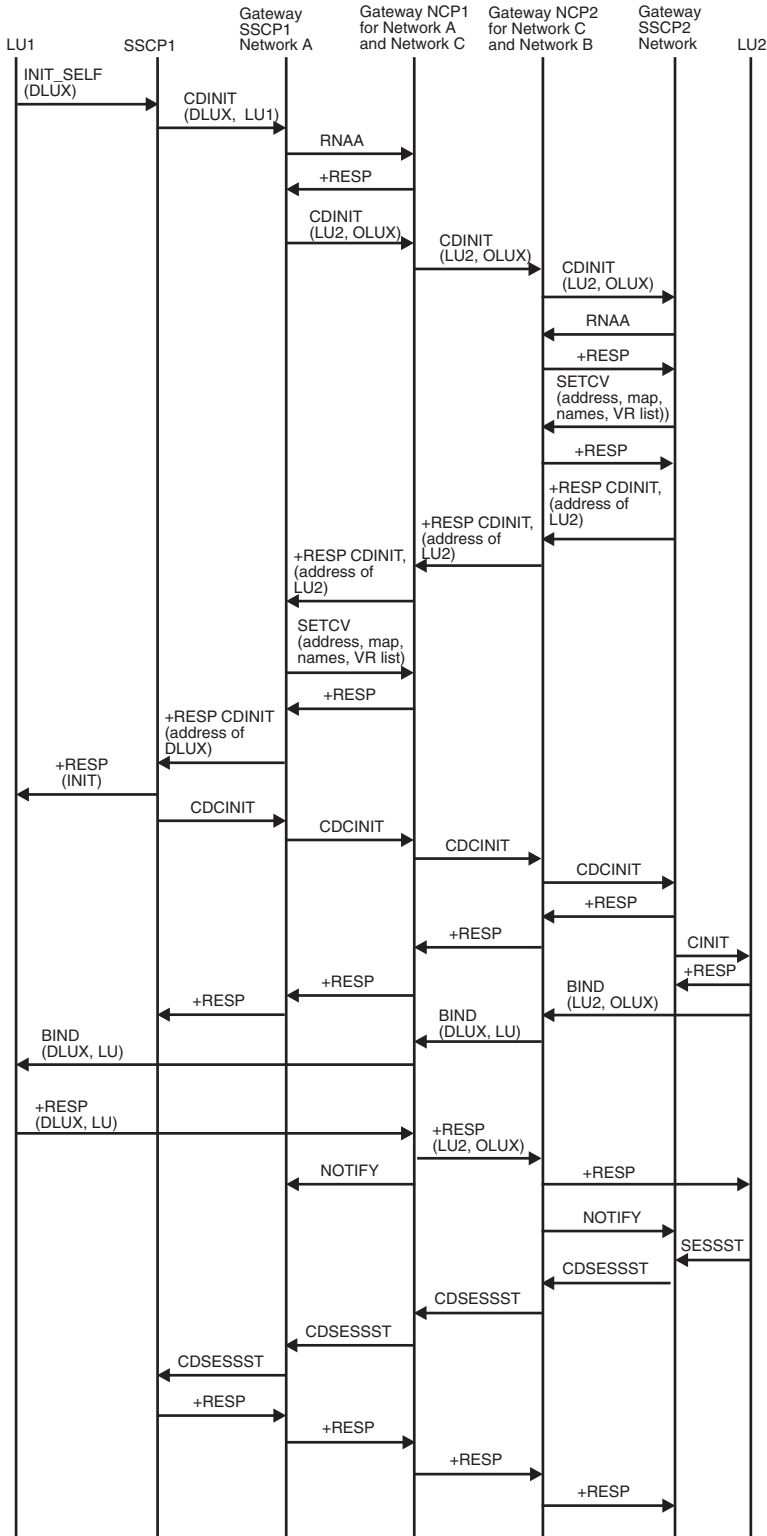


Figure 110. Multiple gateway VTAMs and back-to-back gateway NCPs

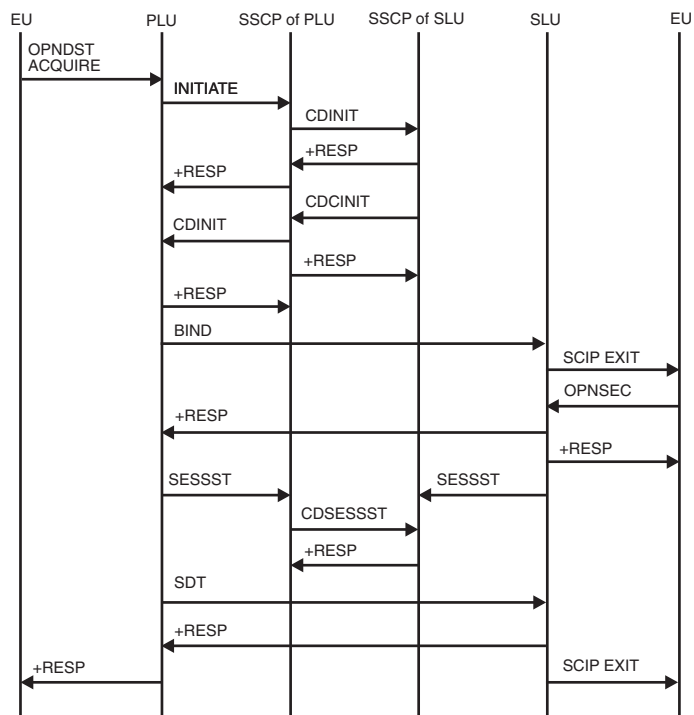


Figure 111. Primary logical unit initiate, OPNDST ACQUIRE

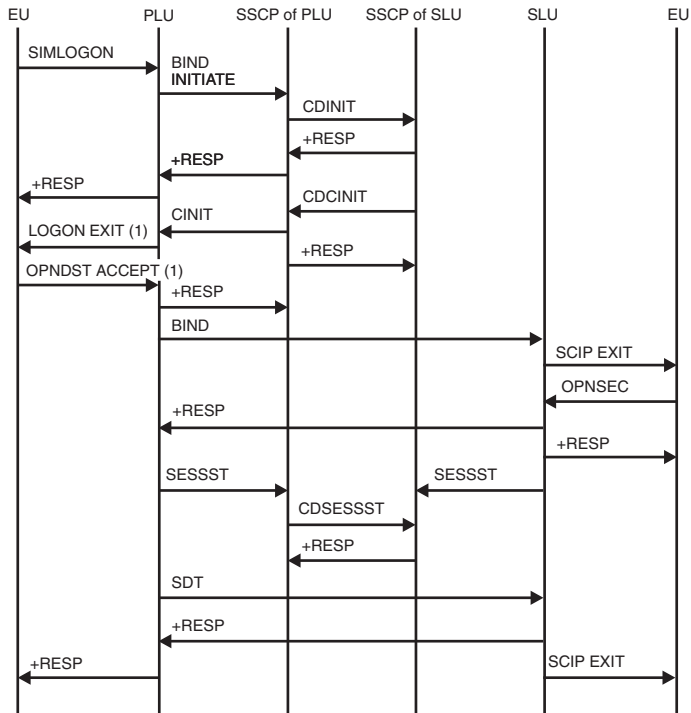


Figure 112. Primary logical unit initiate, SIMLOGON

1. LOGON EXIT and OPNDST ACCEPT flow only when the PLU is associated with an application program. It does not appear in the flow if the PLU is a device-type logical unit. This is true for many following flows with LOGON EXIT and OPNDST ACCEPT.

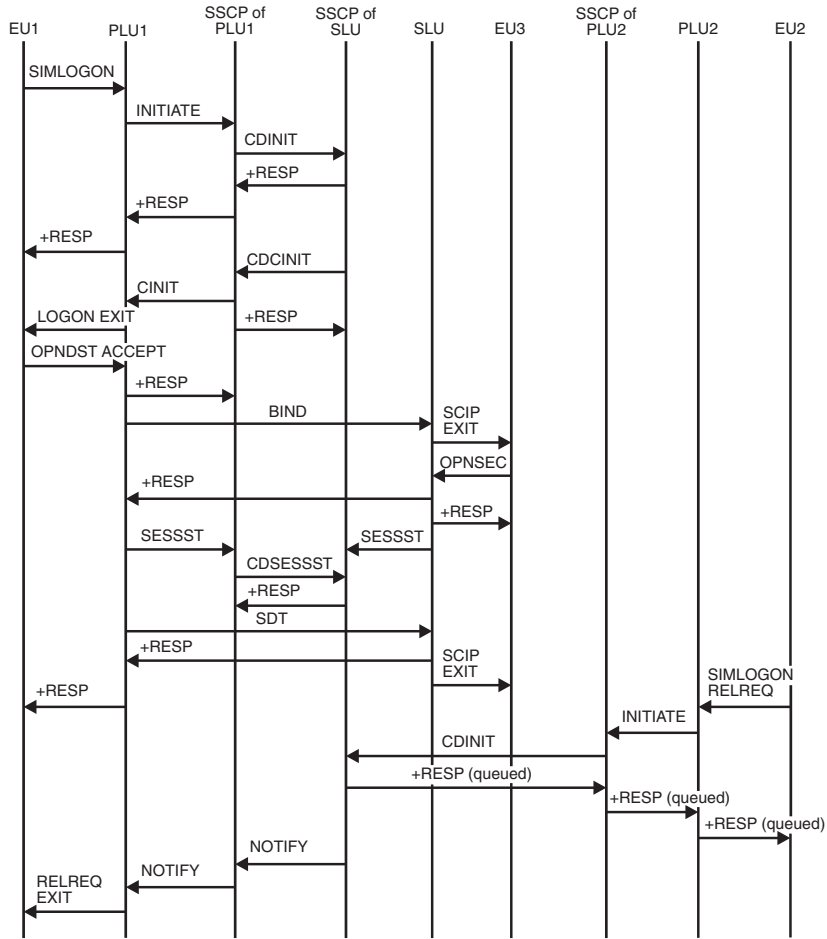


Figure 113. Primary logical unit initiate, SIMLOGON(RELREQ)

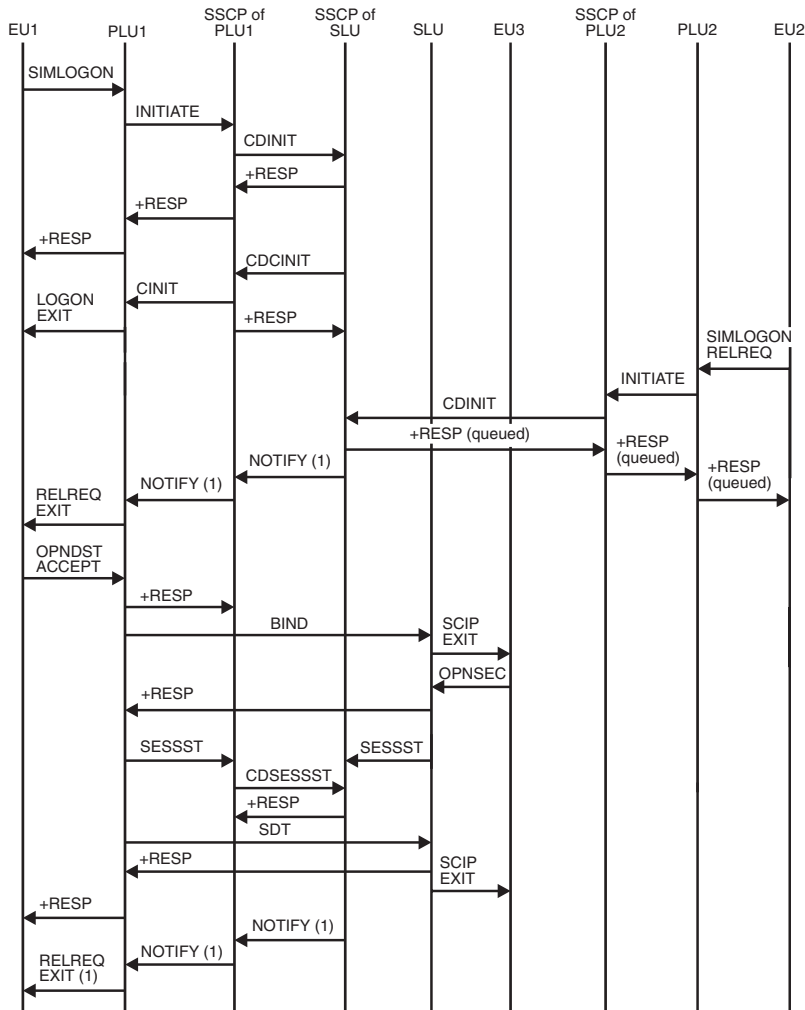


Figure 114. Primary logical unit initiate, SIMLOGON(RELREQ): Session is pending active or already in progress

1. If the session is pending active, the NOTIFY request and the RELREQ EXIT request flow after the response to the CDINIT request from the PLU2 SSCP. If the session is already in progress, these requests flow after the response to the OPNDST ACCEPT.

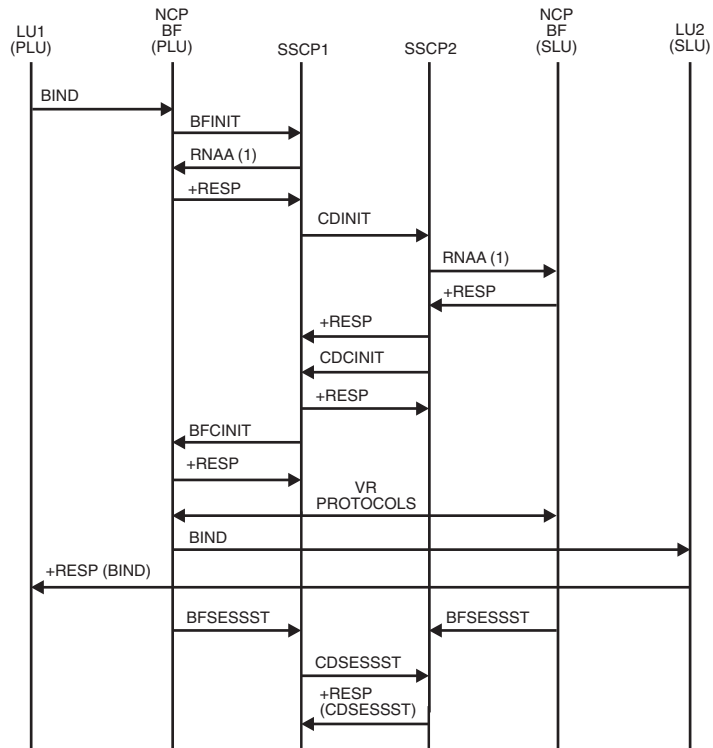


Figure 115. Independent PLU initiating cross-domain session with independent SLU

1. RNAA flows only if the network address is needed.

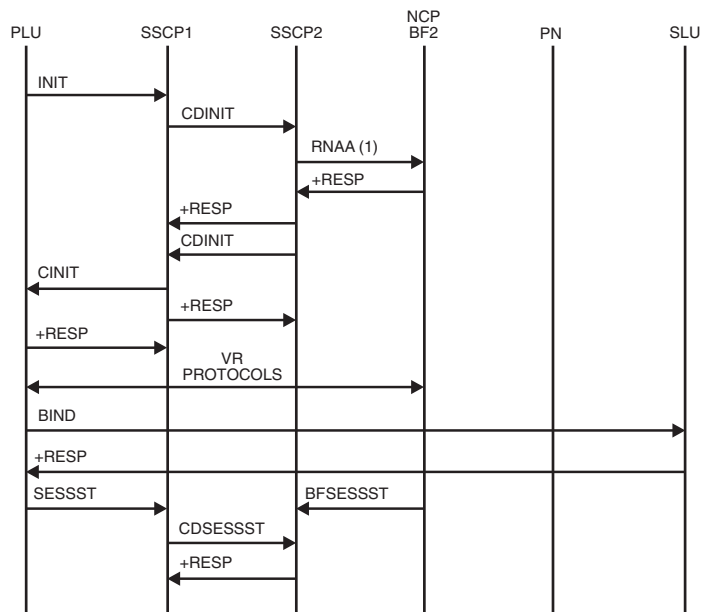


Figure 116. Dependent PLU initiating cross-domain session with independent SLU

1. RNAA flows only if the network address is needed.

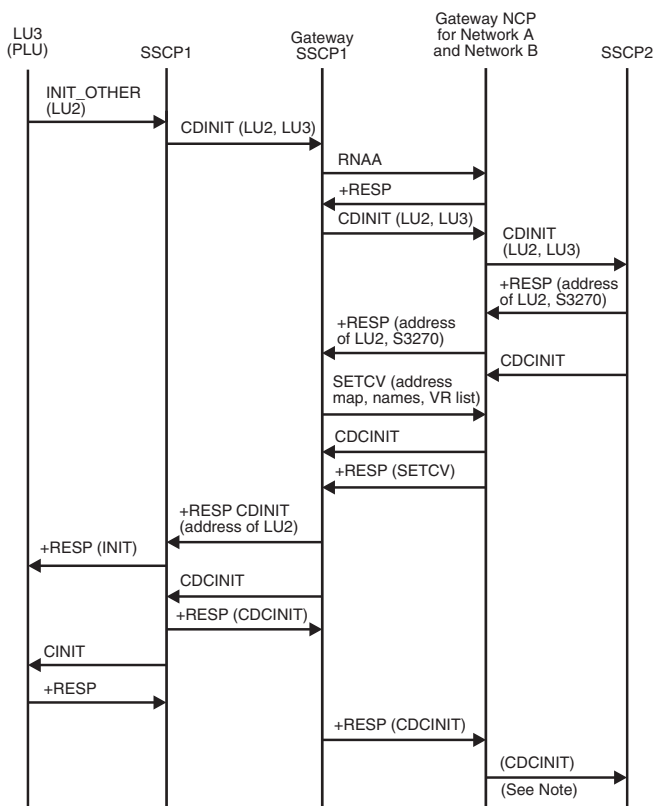


Figure 117. PLU initiating request for single gateway VTAM and single gateway NCP

Note: Session setup continues as in the flow for a SLU-initiated session.

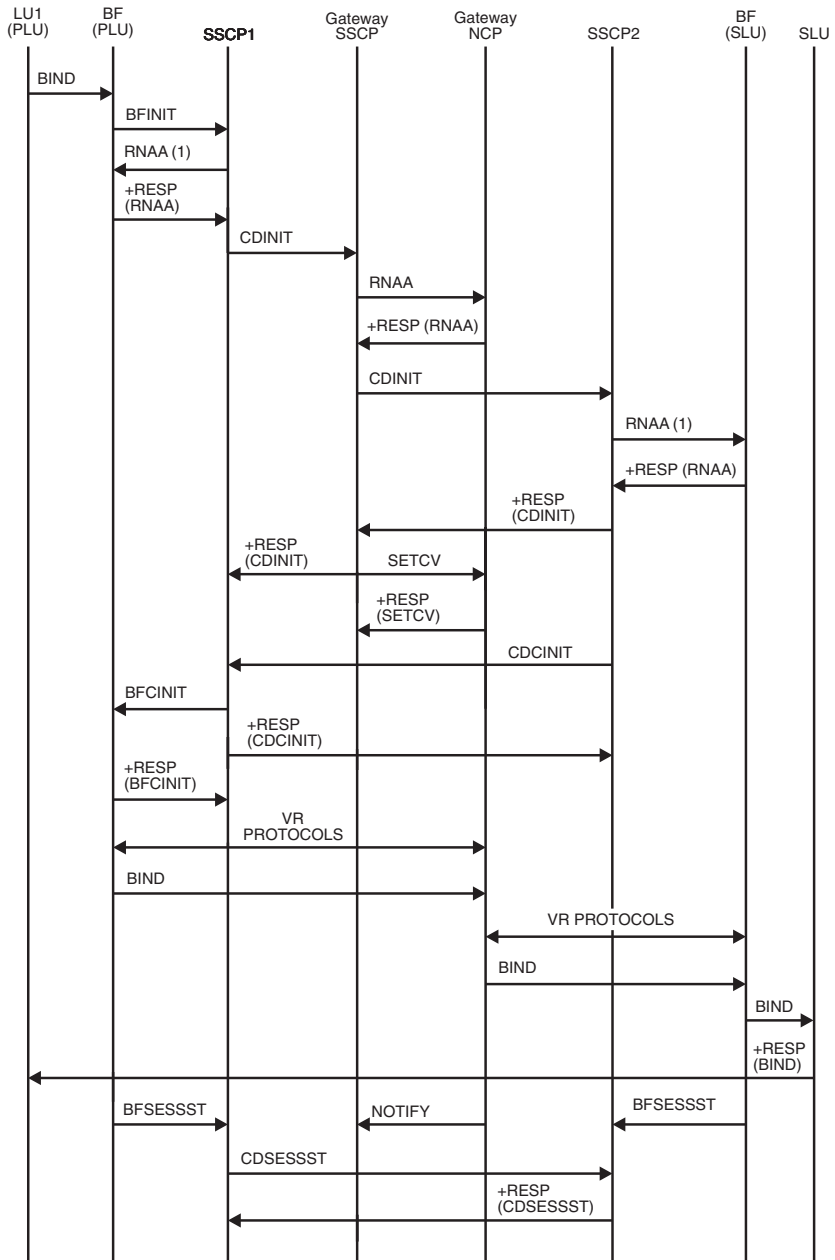


Figure 118. Independent PLU requesting session with independent SLU through single gateway VTAM and single gateway NCP

1. RNAA flows only if the network address is needed.

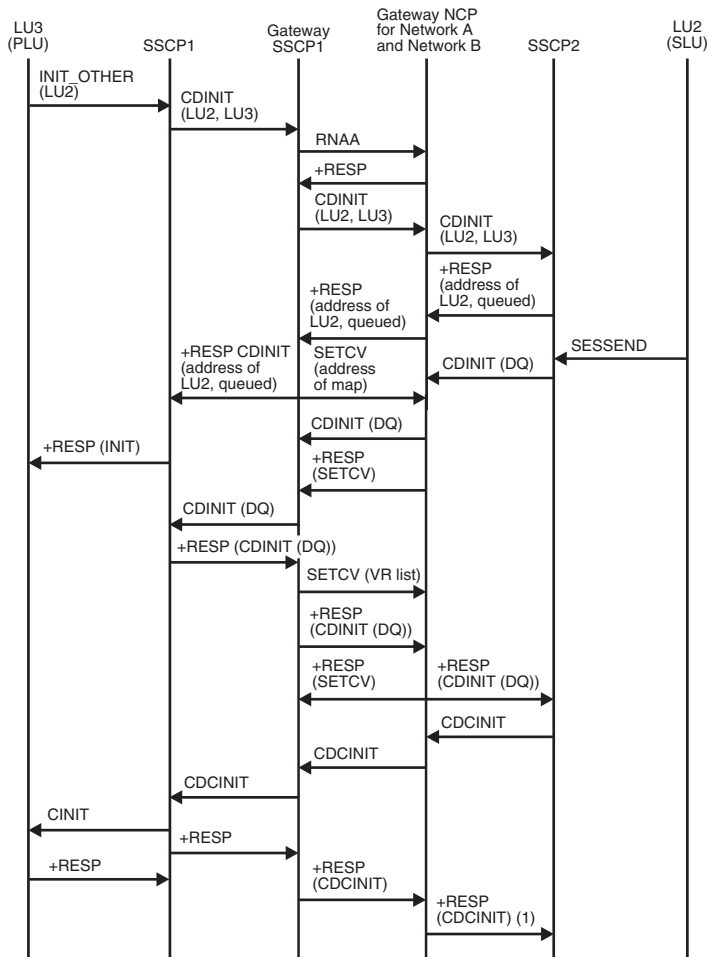


Figure 119. PLU-initiated request setup queued for single gateway NCP and single gateway VTAM

1. Session setup continues as in the flow for an SLU-initiated session.

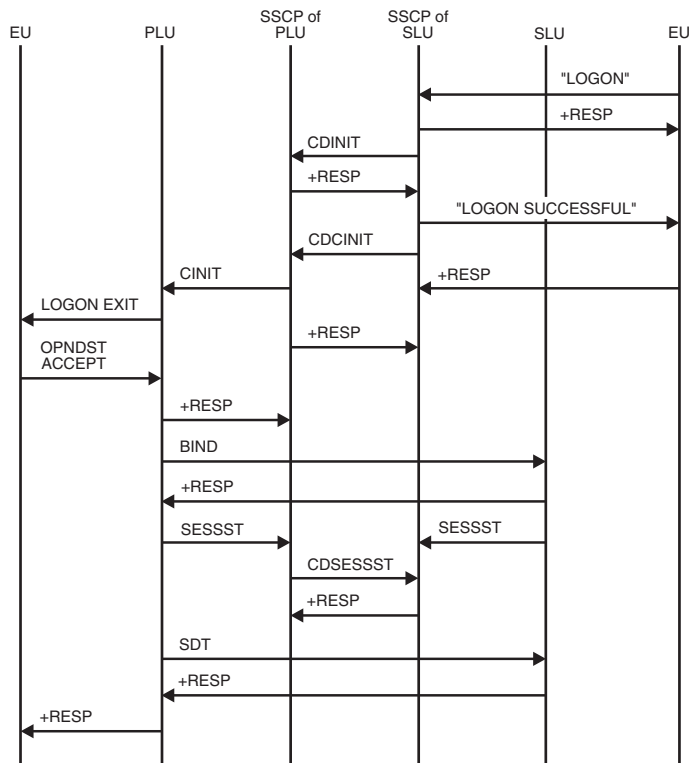


Figure 120. Secondary logical unit initiate (LOGON)

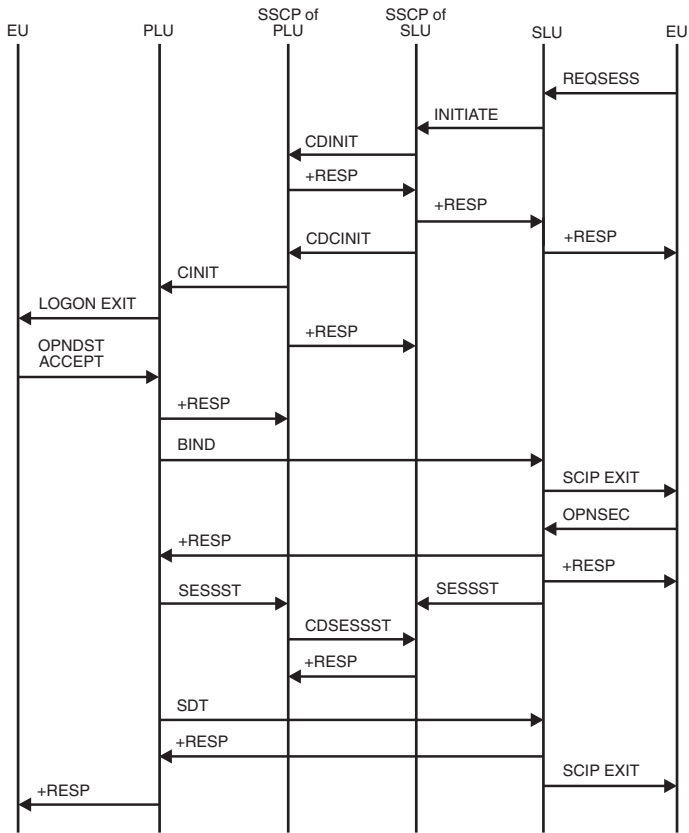


Figure 121. Secondary logical unit initiate (REQSESS)

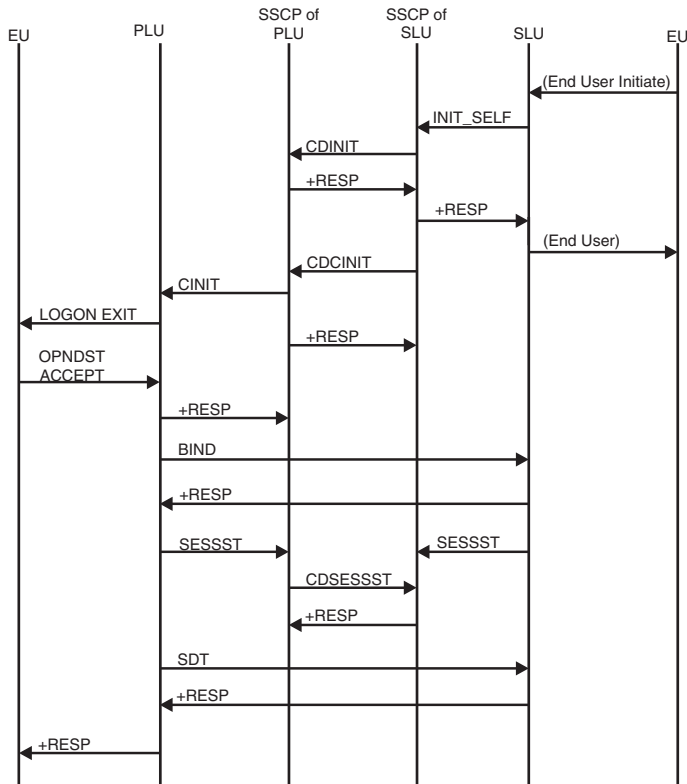


Figure 122. Secondary logical unit initiate (INIT SELF)

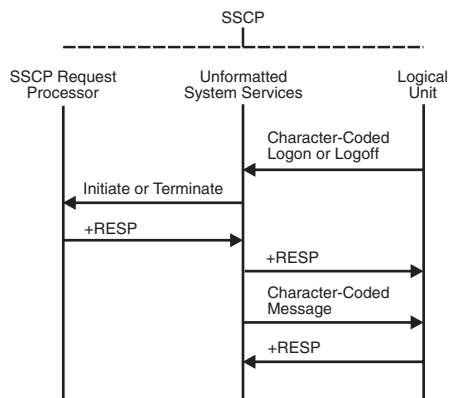


Figure 123. Sending an unformatted request to the SSCP

In this example, the logical unit sends a character-coded logon or logoff to the SSCP. The unformatted system services portion of SSCP converts the logon into a field-formatted Initiate Self or Terminate Self request. The request is then passed to the SSCP request processor.

If the return code indicates an unsuccessful transmission, the unformatted system services portion of SSCP converts the request into a form that can be understood by the terminal logical unit.

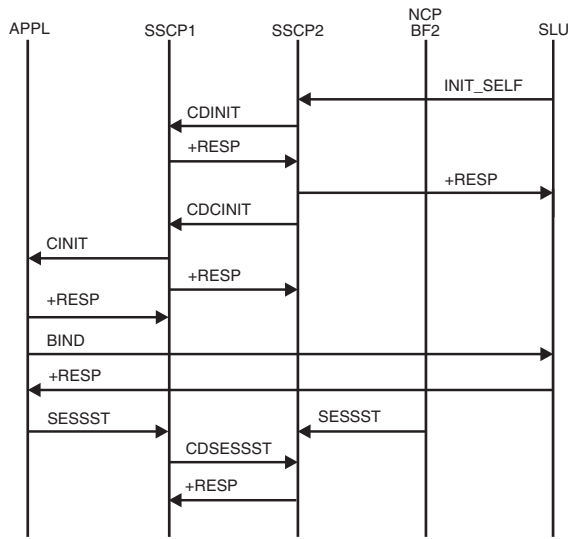


Figure 124. Dependent SLU initiating a cross-domain session with application LU

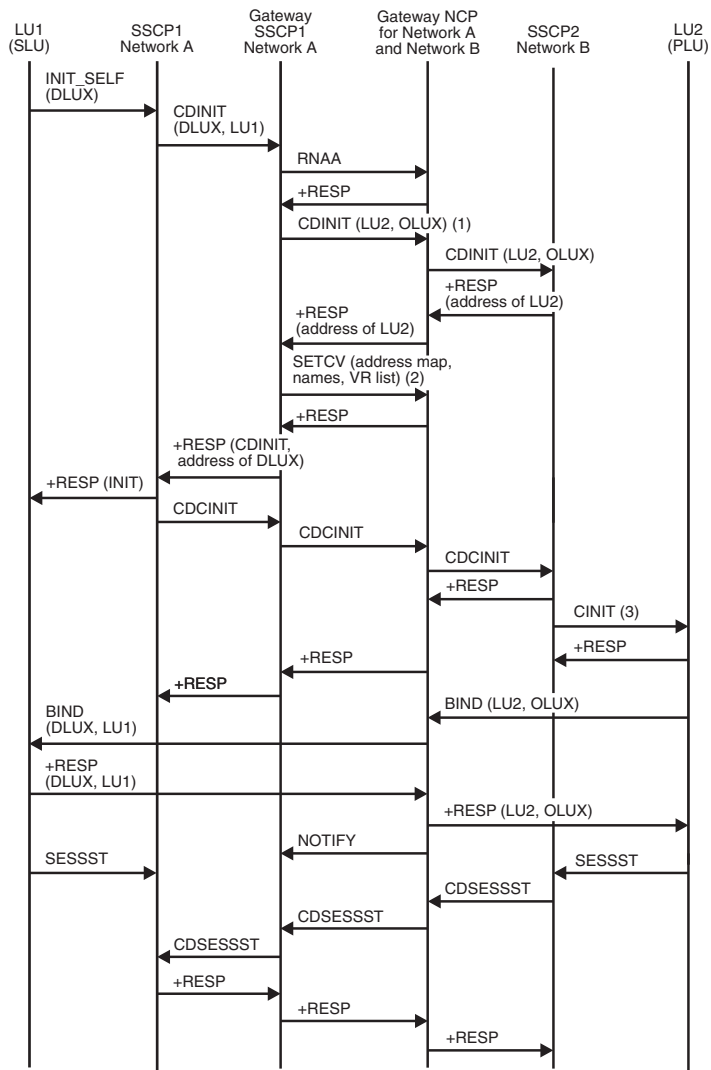


Figure 125. SLU initiating request for single gateway VTAM and single gateway NCP

1. LU1 is initiating a session with DLUX. Using alias name translation, SSCP1 translates DLUX to LU2 and LU1 to OLUX.
2. Names are sent to allow substitution in the BIND.
3. The CINIT drives the logon exit.

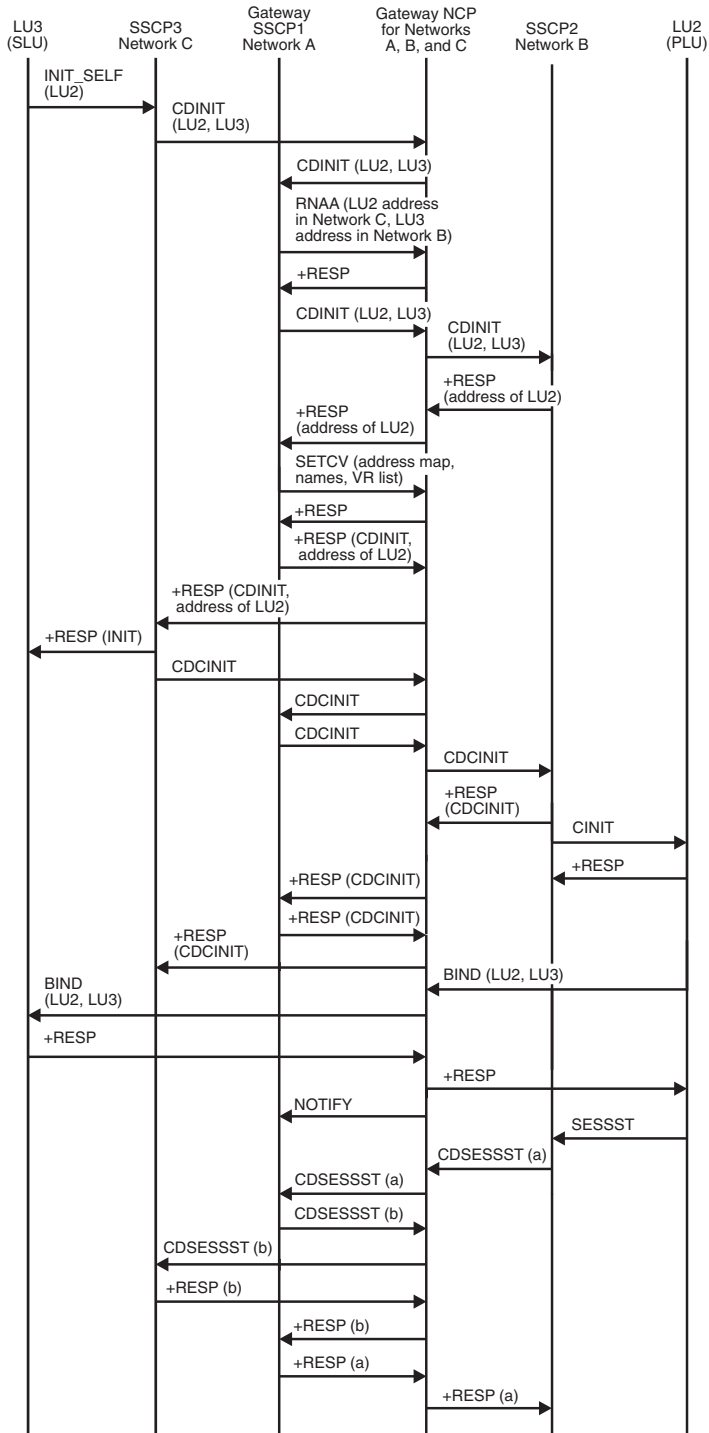


Figure 126. SLU initiating request for single gateway connecting three or more networks

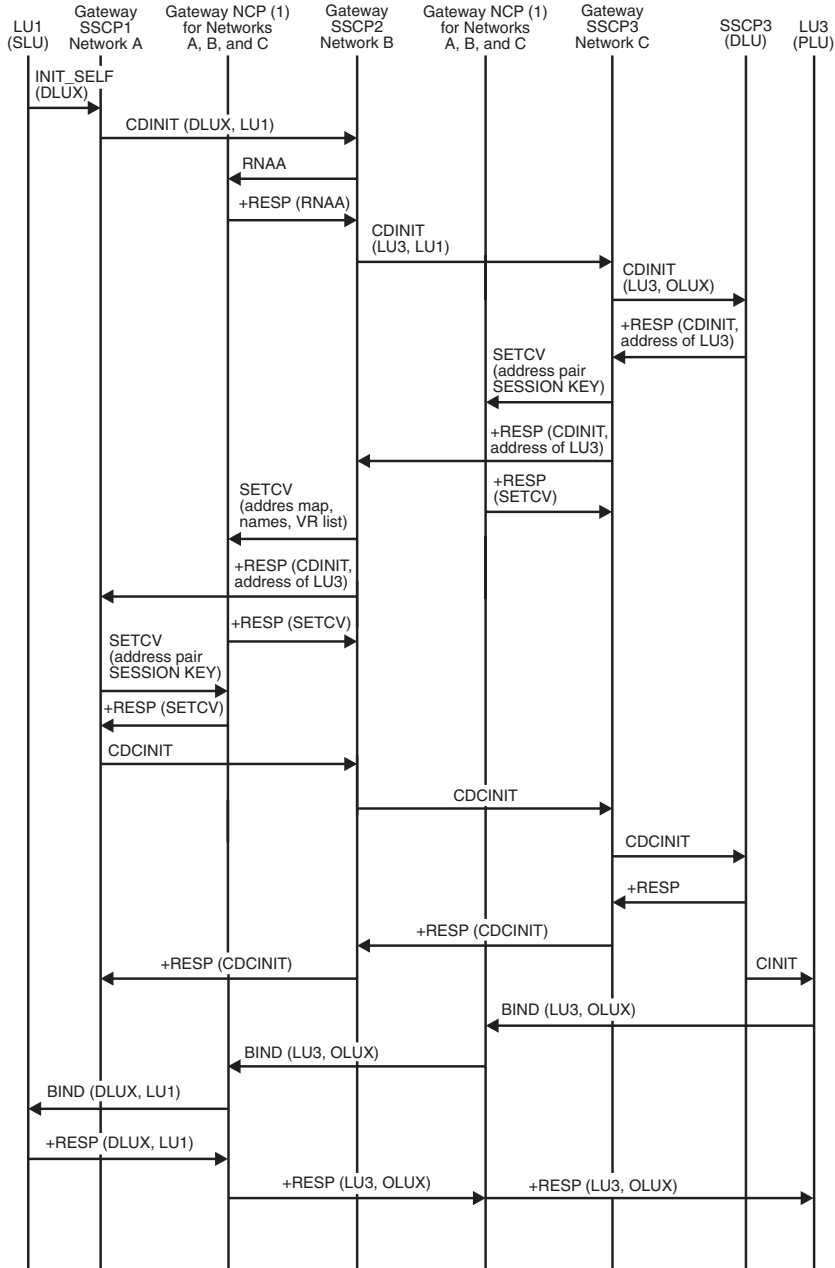


Figure 127. SLU initiating request for predesignated control of gateway NCP by middle host (part 1 of 2)

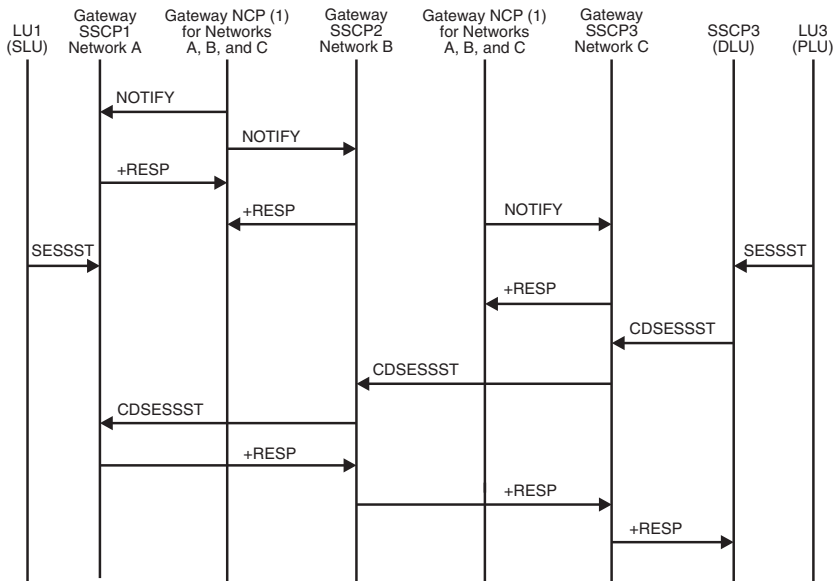


Figure 128. SLU initiating request for predesignated control of gateway NCP by middle host (part 2 of 2)

To simplify the flow, the gateway NCP is shown twice in this flow.

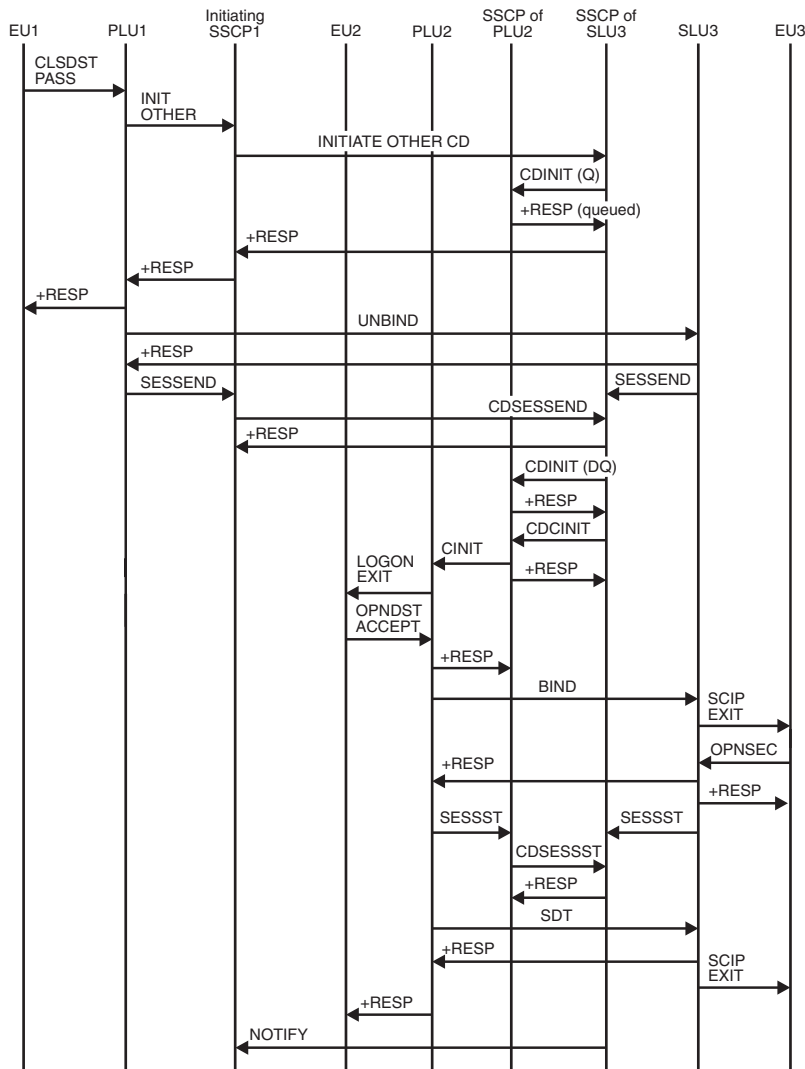


Figure 129. Third party initiating CLSDST PASS

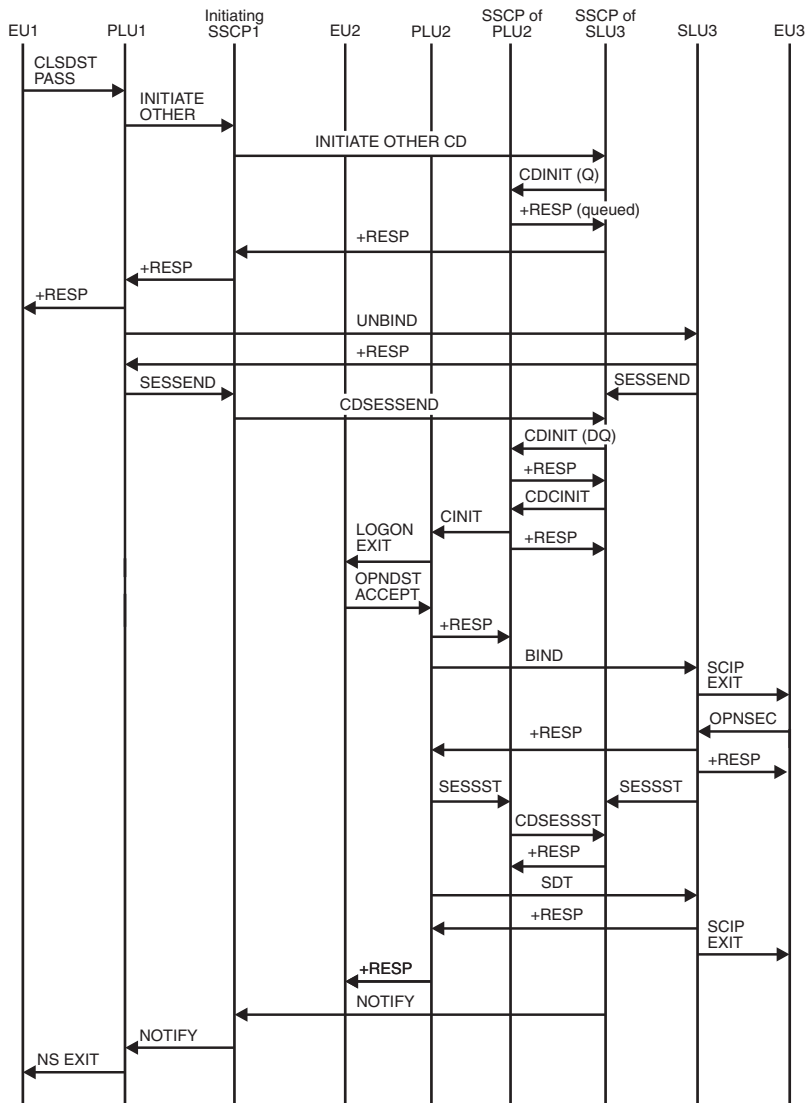


Figure 130. Third party initiating CLSDST PASS with NOTIFY

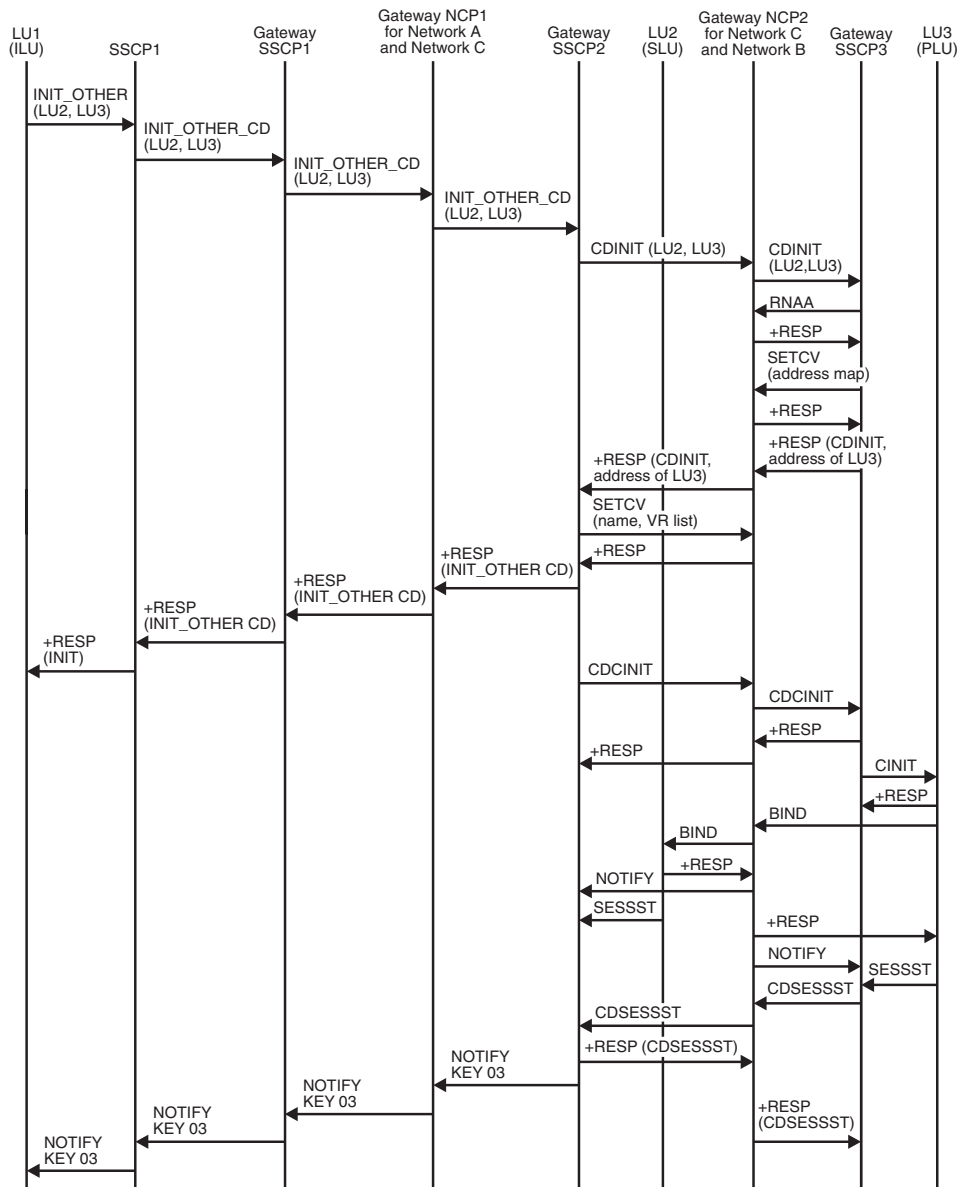


Figure 131. Third party initiating request spanning three networks

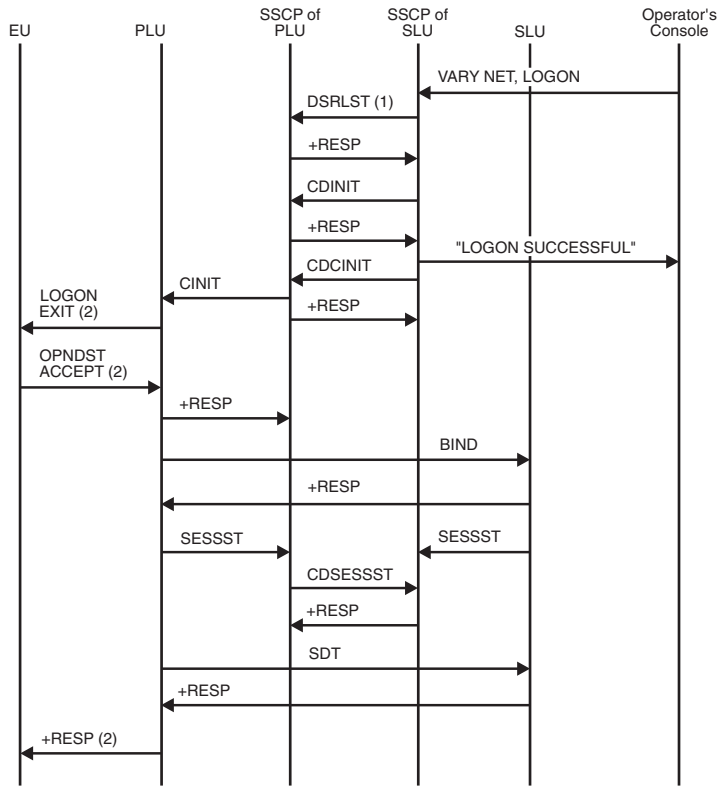


Figure 132. Initiating session using VARY NET, LOGON or LOGAPPL

1. Optional; occurs only when SLU is a dial device.
2. This applies only when the PLU is associated with an application program. It does not appear in the flow if the PLU is a device-type logical unit.

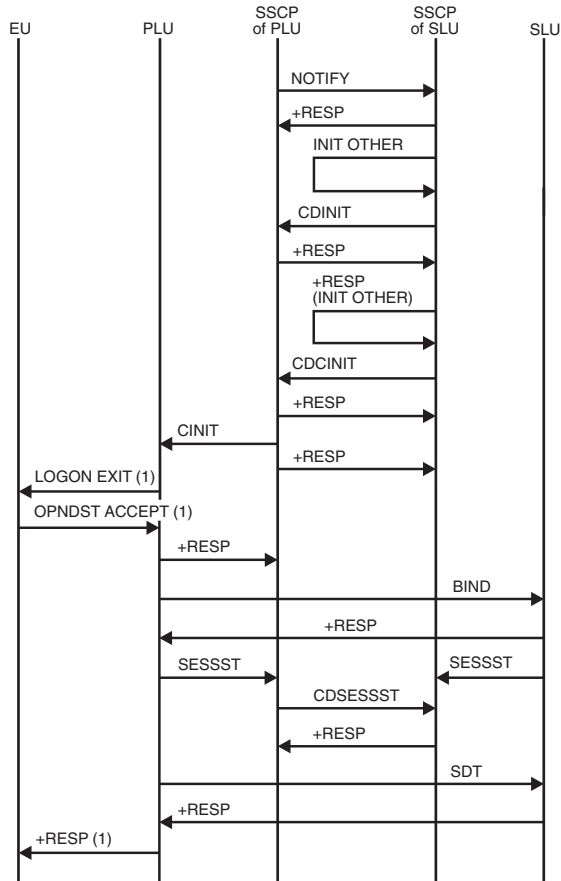


Figure 133. Notification of PLU availability for autologon

1. This applies only when the PLU is associated with an application program. It does not appear in the flow if the PLU is a device-type logical unit.

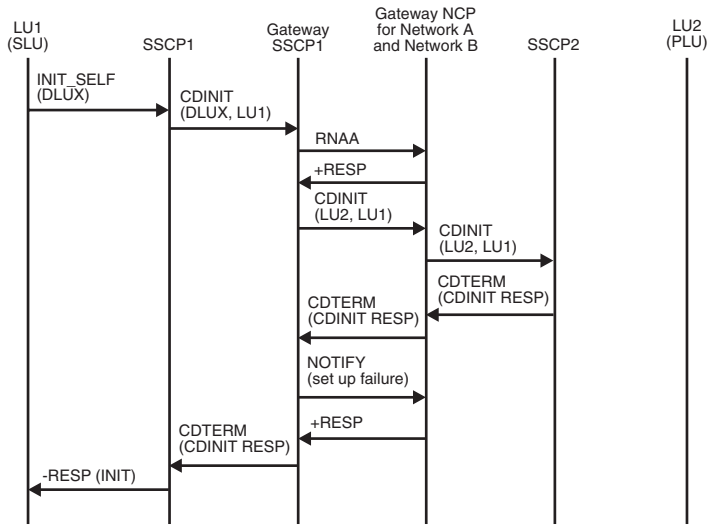


Figure 134. Failure (CDINIT rejection) of session initiated by an SLU for single gateway VTAM and single gateway NCP

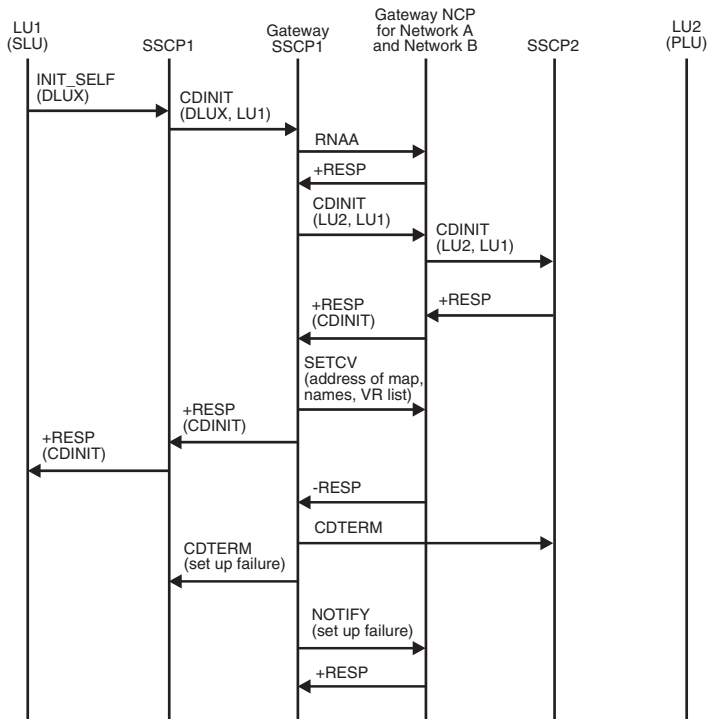


Figure 135. Failure (SETCV failure) of session initiation by an SLU for single gateway VTAM and single gateway NCP

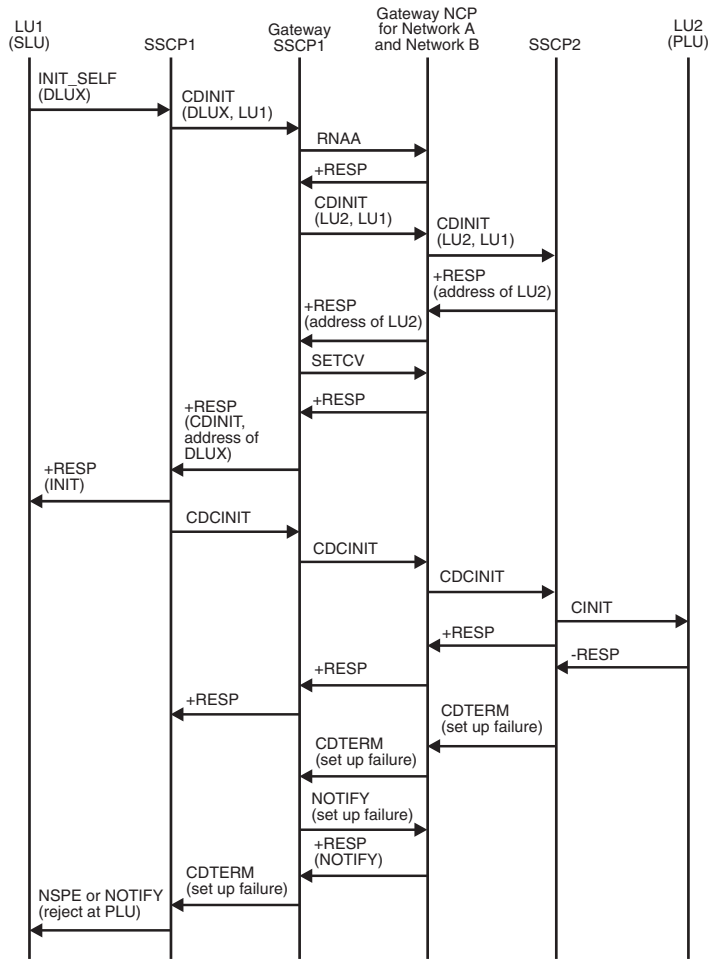


Figure 136. Failure (CINIT rejection) of setup procedure initiated by an SLU for single gateway VTAM and single gateway NCP

Deactivation and session termination flows

Figure 137 on page 494 through Figure 162 on page 511 show the flow of requests and responses between the SSCP and logical and physical units to deactivate resources and end sessions.

Index of deactivation and session termination flows

Table 37 lists the deactivation and session termination flows that are illustrated here.

Table 37. Index of deactivation and session termination flows

Flow	Page
CLOSE ACB processing	Figure 152 on page 504
Deactivating an application program	Figure 153 on page 505
Deactivating a CDRM	

Table 37. Index of deactivation and session termination flows (continued)

Flow	Page
Forced	Figure 156 on page 507
Forced, without affecting active sessions	Figure 158 on page 508
Forced or immediate, VTAM releases before V3R4.1	Figure 159 on page 509
Immediate	Figure 155 on page 507
Immediate, without affecting active sessions	Figure 157 on page 508
Normal	Figure 154 on page 506
Deactivating a logical unit (LU), single network	
Forced	Figure 138 on page 494
Immediate	Figure 137 on page 494
VARY NET,TERM Cleanup	Figure 150 on page 503
VARY NET,TERM Unconditional	Figure 149 on page 502
With Giveback	Figure 139 on page 495
Deactivating a logical unit (LU), multiple networks	
Independent PLU sends BFCLEANUP for independent SLU	Figure 140 on page 495
Independent PLU sends UNBIND for independent SLU	Figure 141 on page 496
PLU sends UNBIND for multiple gateway VTAMs and single gateway NCP	Figure 142 on page 496
PLU sends UNBIND for single gateway VTAM and single gateway NCP	Figure 143 on page 497
SLU requests TERMINATE SELF (CLEANUP) for single gateway VTAM and single gateway NCP	Figure 145 on page 499
SLU requests TERMINATE SELF for multiple gateway VTAMs and back-to-back gateway NCPs	Figure 144 on page 498
SLU requests TERMINATE SELF for single gateway VTAM and single gateway NCP	Figure 146 on page 500
Type 2.1 nodes, active termination	Figure 147 on page 501
Deactivating a physical unit (PU) acting as an adjacent link station for independent logical unit (LU) sessions	
Queued session, terminating	Figure 151 on page 503
Route failure	
Route failure in intermediate network causes termination of LU-LU sessions	Figure 161 on page 511
Route failure in intermediate network causes termination of SSCP-SSCP sessions	Figure 162 on page 511
SSCP-SSCP session termination causes LU-LU sessions to be broken	Figure 160 on page 510

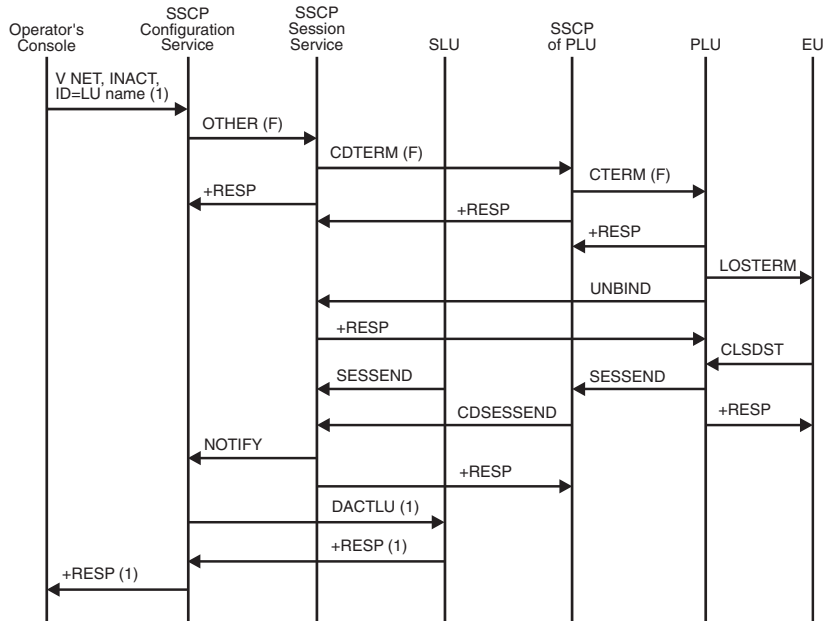


Figure 137. Deactivating a logical unit: Immediate

1. These flow only when the operator deactivates a specific logical unit. For example, they do not flow during immediate deactivation of a CDRM.

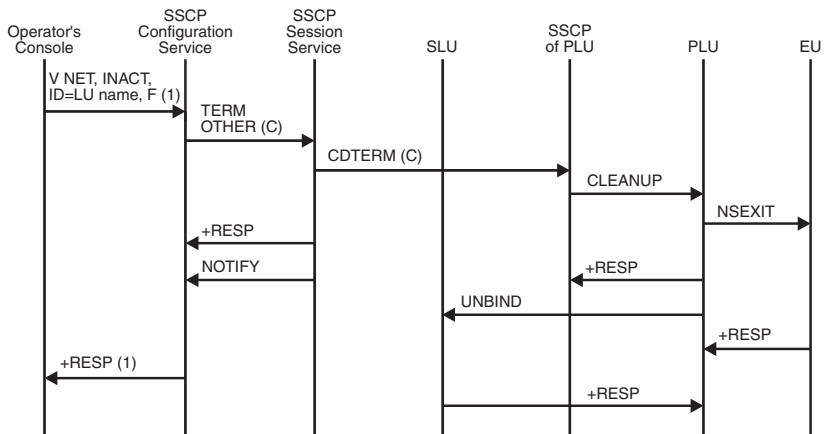


Figure 138. Deactivating a logical unit: Forced

1. These flow only when the operator deactivates a specific logical unit. For example, they do not flow during forced deactivation of a CDRM.

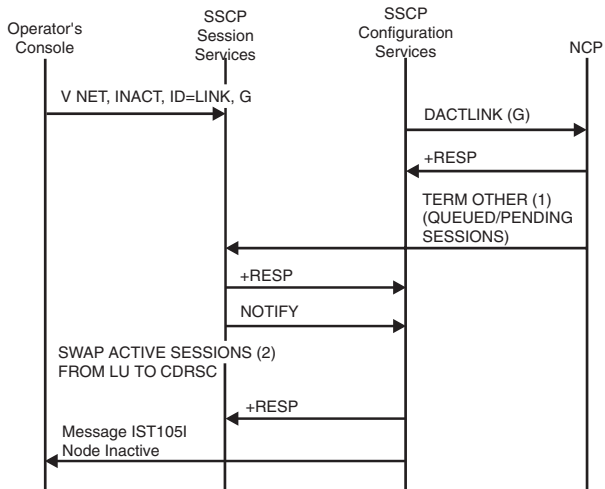


Figure 139. Deactivating a logical unit with giveback

1. The DACTLINK X'02' terminates only queued and pending LU-LU sessions. Active LU-LU sessions remain active.
2. After session services transfers SIBs of ACTIVE logical units to the CDRSC, configuration services SRTADDs the CDRSCs as real resources, and the logical units are ADDED as shadow resources. If a CDRSC for a particular logical unit does not exist, a dynamic CDRSC is allocated for the logical unit.

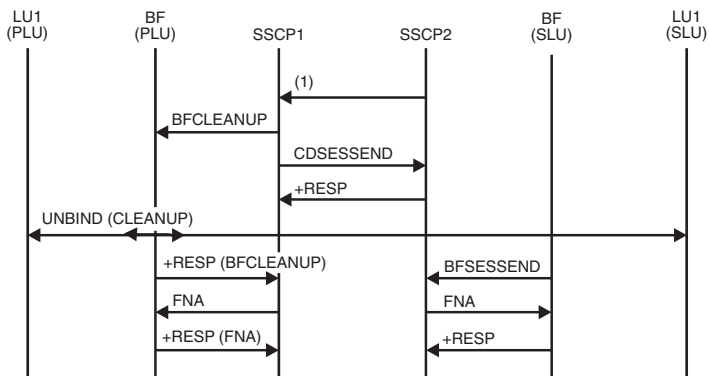


Figure 140. Independent primary logical unit (PLU) sends BFCLEANUP for cross-domain LU-LU session with independent secondary logical unit (SLU)

- BFCLEANUP can be sent by the SSCP(PLU) for several reasons, including the following conditions:
 - A network operator at the SSCP(PLU) issues a VARY NET,TERM,UNCOND, generating an internal TERM-OTHER(forced).
 - A network operator at either SSCP issues a VARY NET,INACT,ID=cdrm, deactivating all cross-domain sessions between the SSCPs.

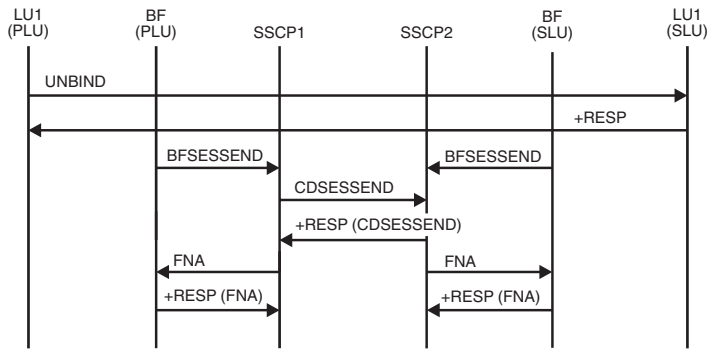


Figure 141. Independent primary logical unit (PLU) sends UNBIND for cross-domain LU-LU session with independent secondary logical unit (SLU)

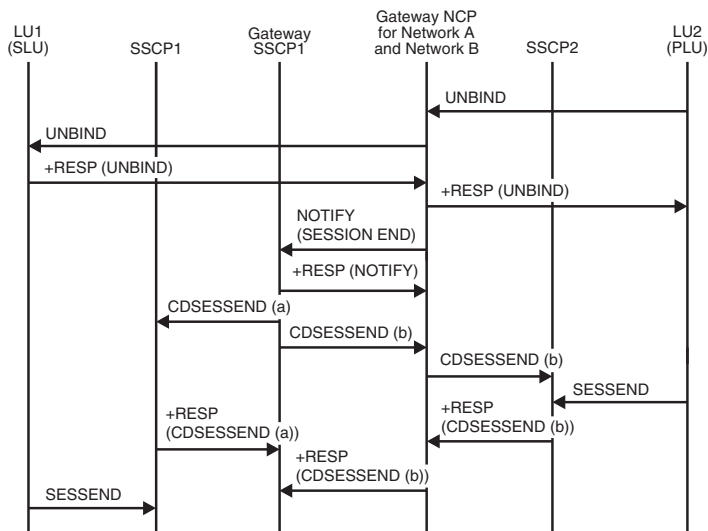


Figure 142. Primary logical unit (PLU) sends UNBIND for multiple gateway VTAMs and single gateway NCP

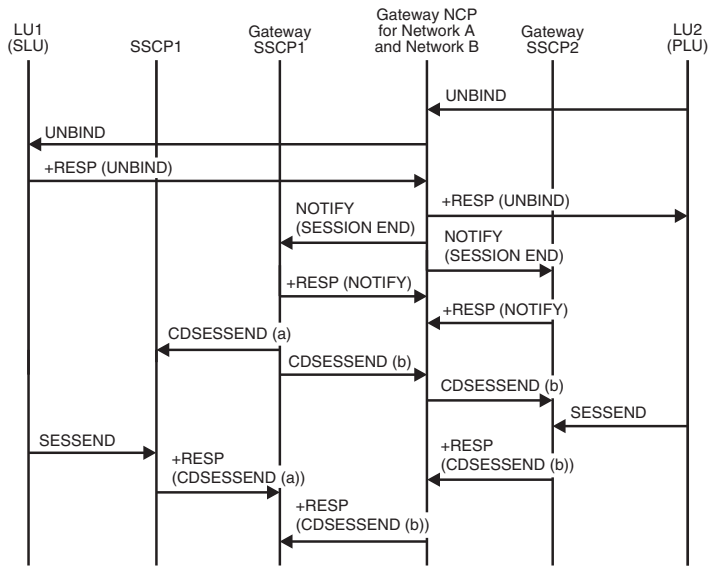


Figure 143. Primary logical unit (PLU) sends UNBIND for single gateway VTAM and single gateway NCP

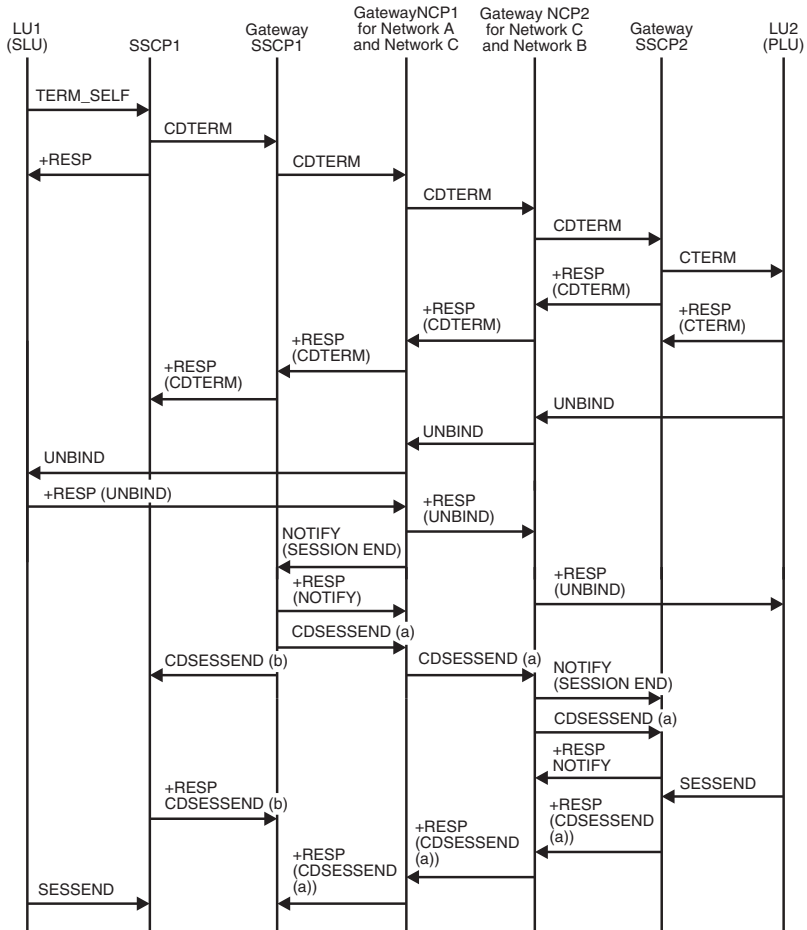


Figure 144. Secondary logical unit (SLU) requests TERMINATE SELF for multiple gateway VTAMs and back-to-back gateway NCPs

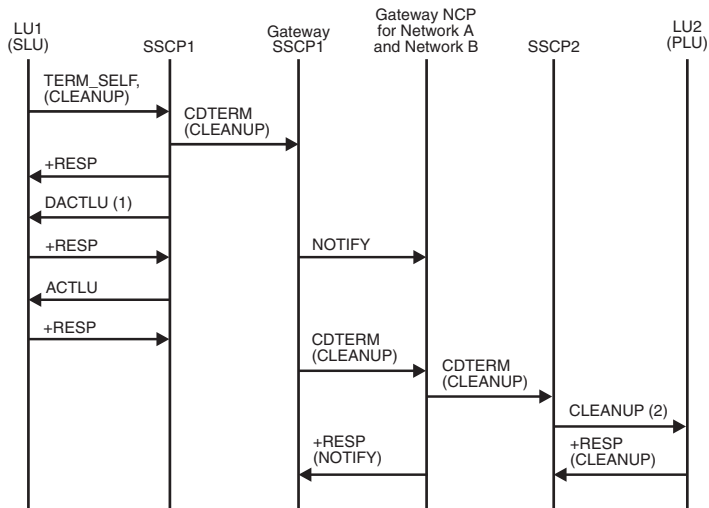


Figure 145. Secondary logical unit (SLU) requests TERMINATE SELF (CLEANUP) for single gateway VTAM and single gateway NCP

Note: The UNBIND can flow from the SLU, the PLU, or the gateway NCP.

1. A DACTLU does not flow to a binary synchronous communication (BSC) terminal.
2. You might receive sense code 081E0003, indicating that cleanup has already occurred.

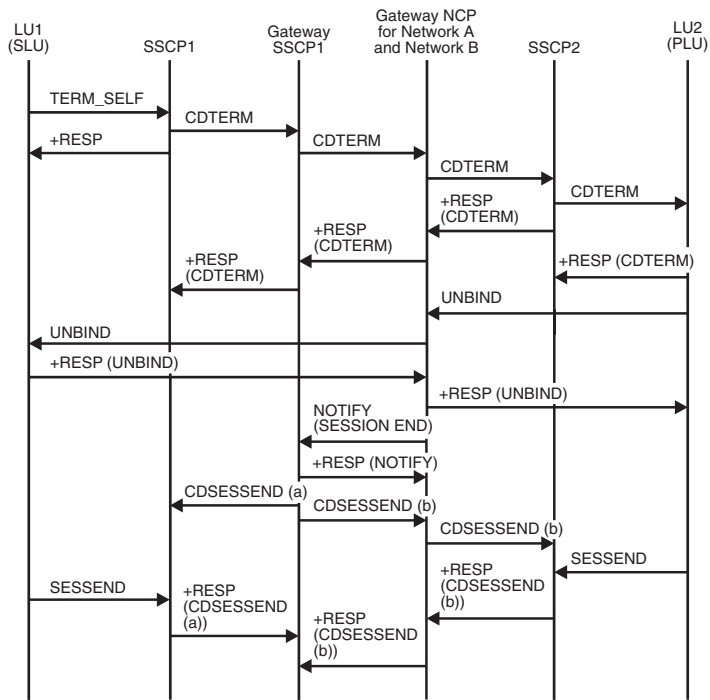


Figure 146. Secondary logical unit (SLU) requests TERMINATE SELF for single gateway VTAM and single gateway NCP

Note: (a) and (b) are used here to differentiate between similar request units.

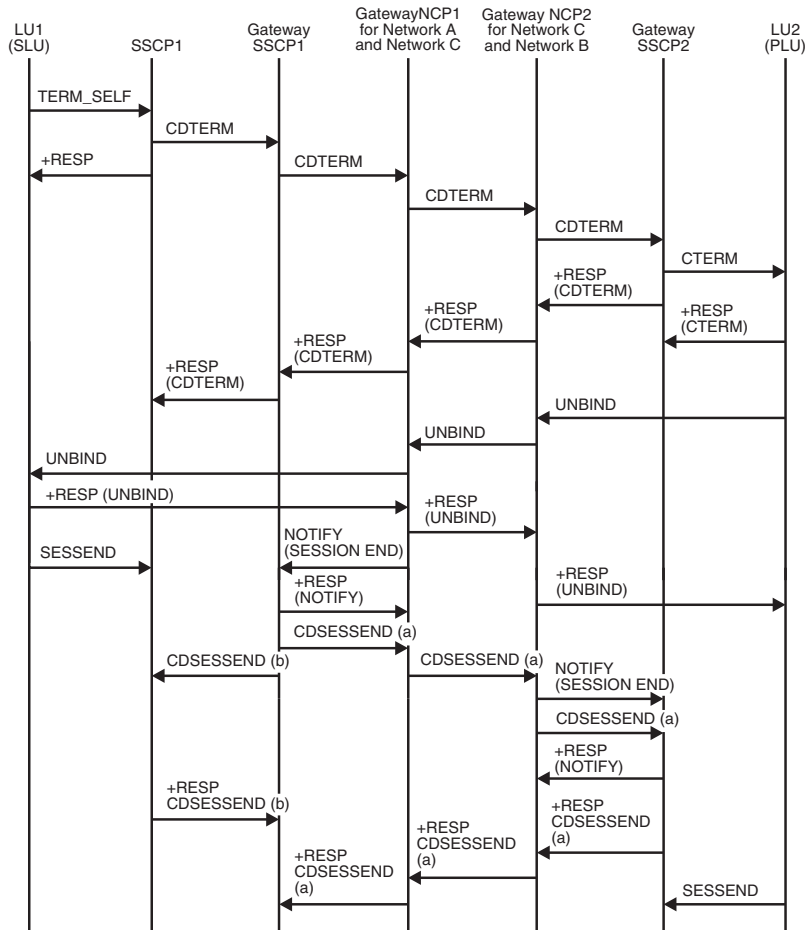


Figure 147. Active session termination of type 2.1 nodes

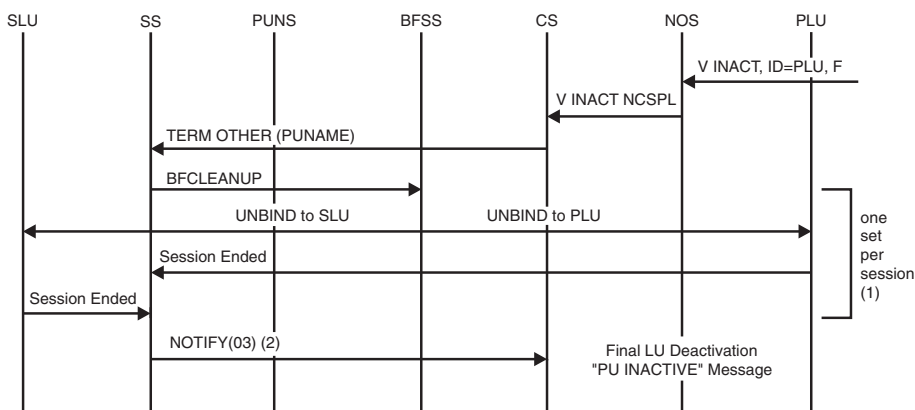


Figure 148. Deactivating a PU acting as an adjacent link station for independent LU sessions

1. Examines PLU and SLU chain and sends (BF)CLEANUP for each ILU session found.

- When all Session Ended (or BFSESEND) signals are received, NOTIFY is sent to CS for PU.

In this example, PU1 is a fictitious adjacent link station. When the PU is deactivated, configuration services sends a TERMINATE containing the PU name to session services. Session services examines the adjacent link station's SIB chains and sends CLEANUP to terminate the sessions. When all sessions are down, NOTIFY flows to configuration services so the final deactivation can occur.

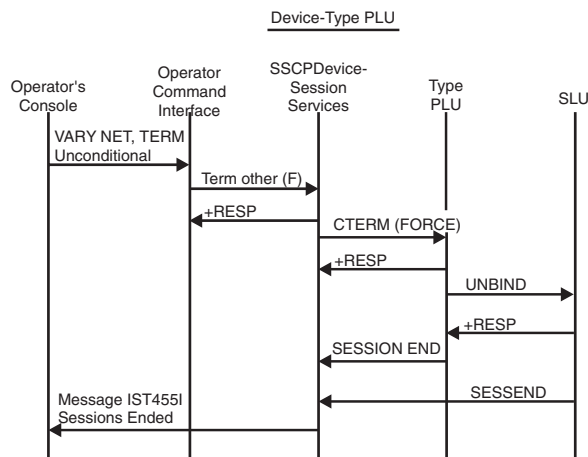
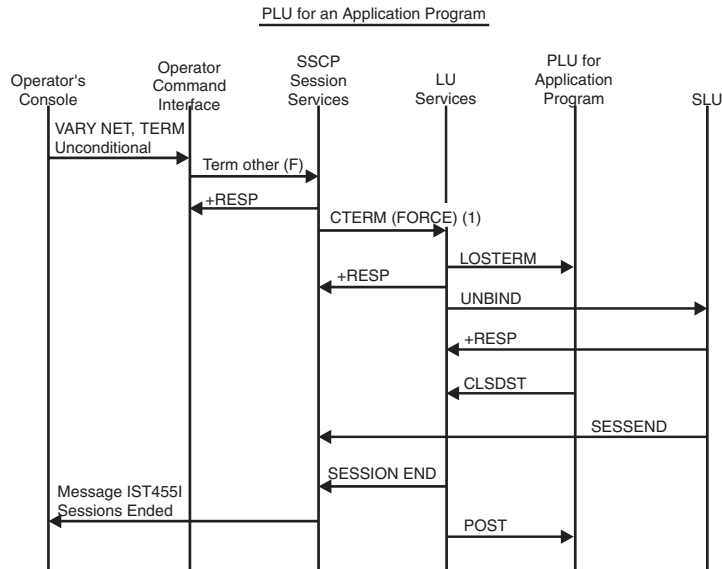


Figure 149. Deactivating sessions or LUs using VARY NET, TERM unconditional

- If the LOSTERM exit routine is already scheduled with a reason code 32 caused by a CTERM (orderly) request that was received before, the CTERM (force) request is upgraded to a CLEANUP RU and VTAM drives an NSEXIT exit routine.

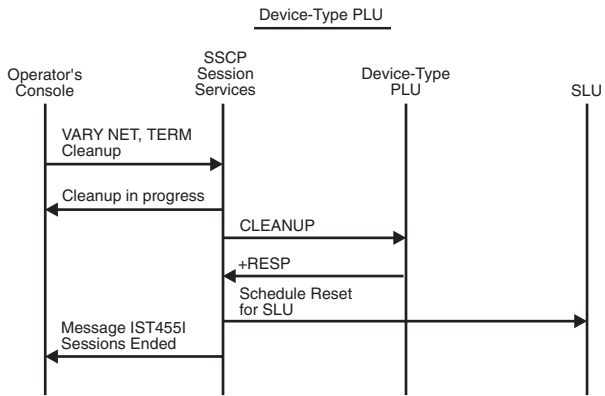
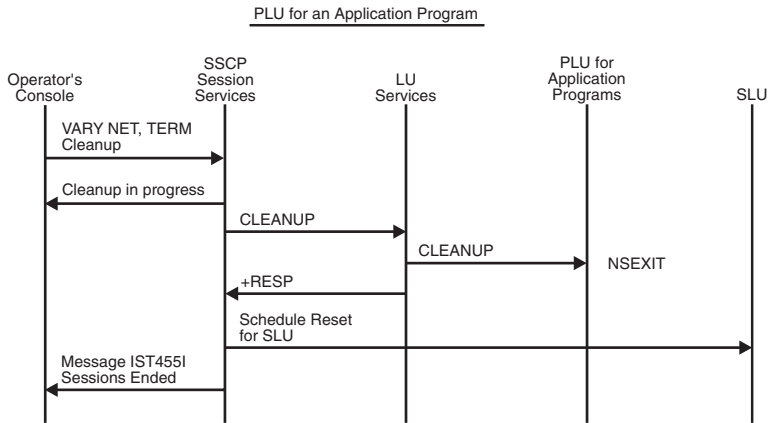


Figure 150. Deactivating sessions or LUs using VARY NET,TERM cleanup

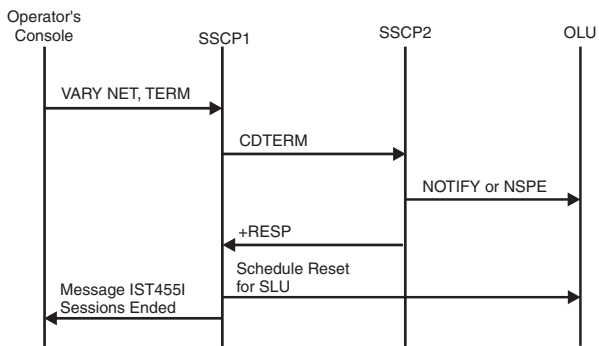


Figure 151. Terminating a queued session

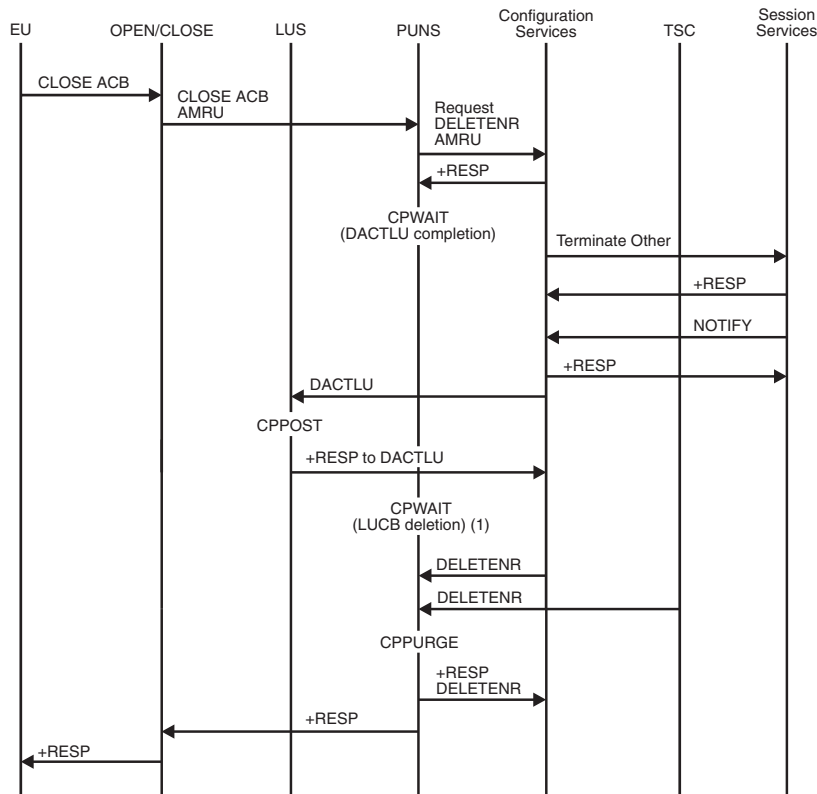


Figure 152. CLOSE ACB processing

1. PUNS cannot send a response to the CLOSE ACB AMRU until DACTLU processing is complete and the LUCB has been deleted. Therefore, after requesting that configuration services deactivate the logical unit, PUNS issues CPWAIT and waits for LUS to post it when the logical unit has been deactivated. After it is posted, PUNS waits to be notified that there are no more active sessions for the application program. PUNS issues CPWAIT and waits for configuration services and TSC to send a request to delete the LUCB. PUNS posts itself when it has processed each of these requests and sends a response to configuration services to notify it that the LUCB has been deleted. After sending this response, PUNS sends a response to the CLOSE ACB AMRU.

For the open ACB flow, see Figure 102 on page 462.

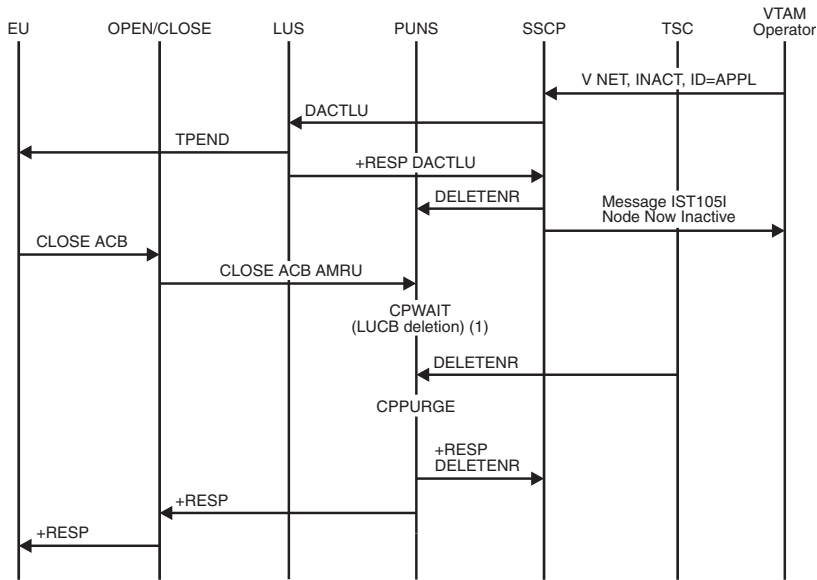


Figure 153. Deactivating an application program

1. PUNS cannot send a response to the CLOSE ACB AMRU until there are no more active sessions for the application program. Therefore, after the SSCP sends a request to delete the LUCB, PUNS waits for OPEN/CLOSE to send it a CLOSE ACB AMRU. When it has received this request, PUNS issues CPWAIT and waits for TSC to send a request to delete the LUCB. PUNS posts itself when it has deleted the LUCB and sends a response to configuration services. After sending this response, PUNS sends a response to the CLOSE ACB AMRU.

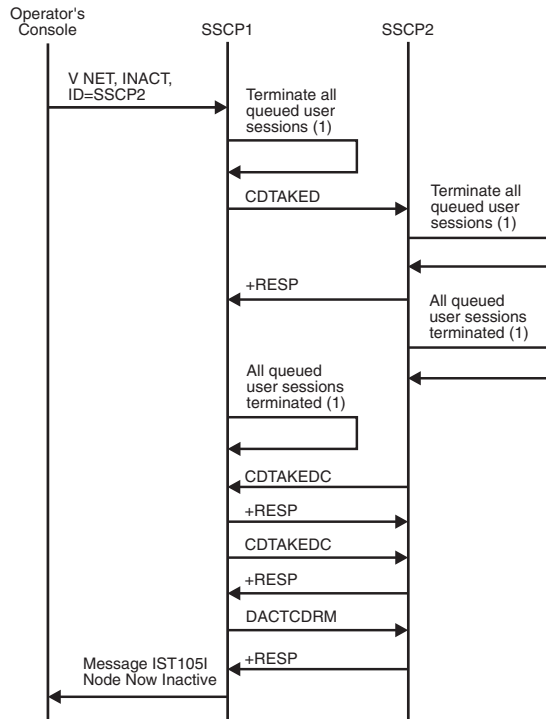


Figure 154. Deactivating a CDRM: Normal

1. See Figure 151 on page 503 for the RUs that flow for termination of a queued session.

Note: When the CDRM is deactivated, immediate processing takes place. See Figure 155 on page 507 for the RUs that flow for immediate deactivation of a CDRM.

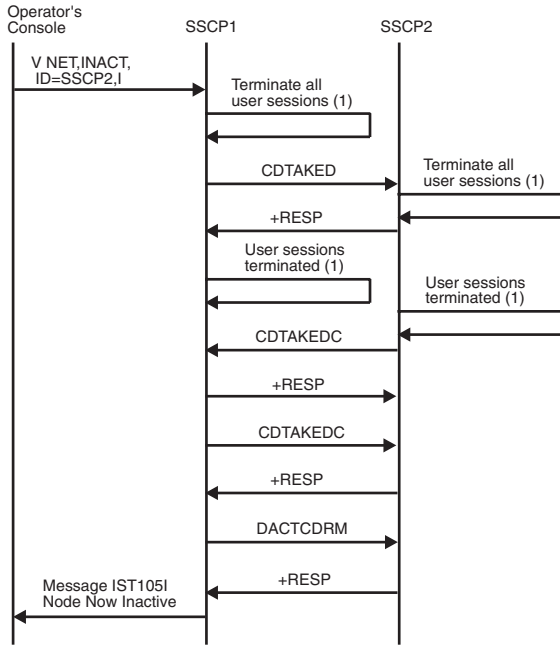


Figure 155. Deactivating a CDRM: Immediate

1. The logical unit will not be deactivated. See Figure 150 on page 503 for the RUs that flow for immediate deactivation of a logical unit.

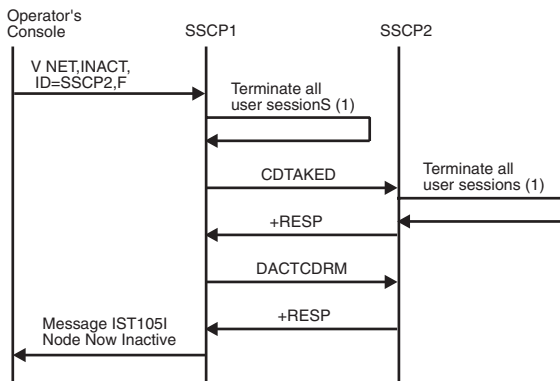


Figure 156. Deactivating a CDRM: Forced

1. The logical unit will not be deactivated. See Figure 150 on page 503 for the RUs that flow for forced deactivation of a logical unit.

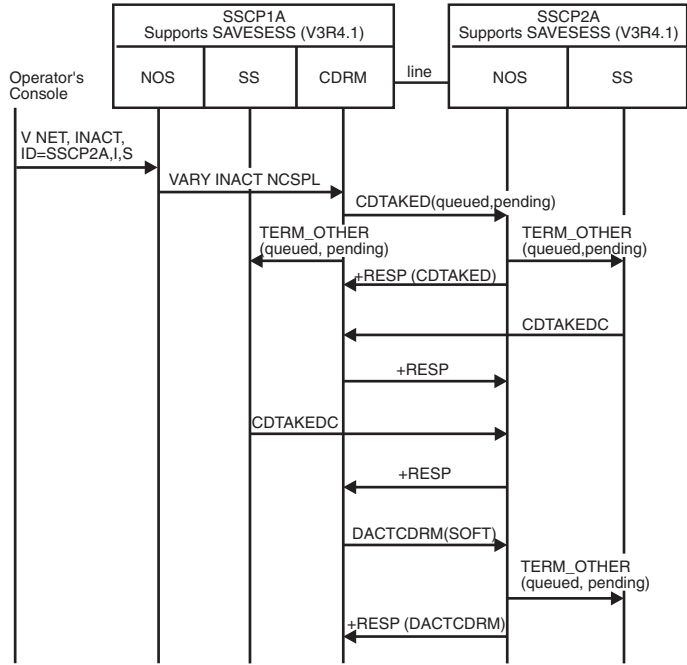


Figure 157. Deactivating a CDRM without affecting active sessions: Immediate

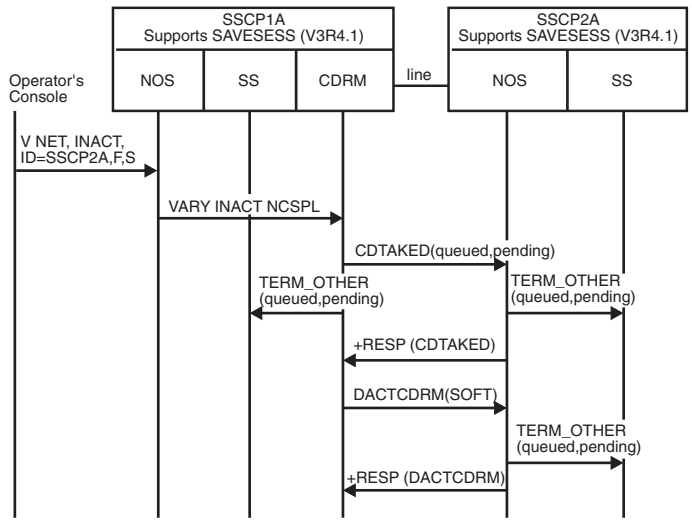


Figure 158. Deactivating a CDRM without affecting active sessions: Forced

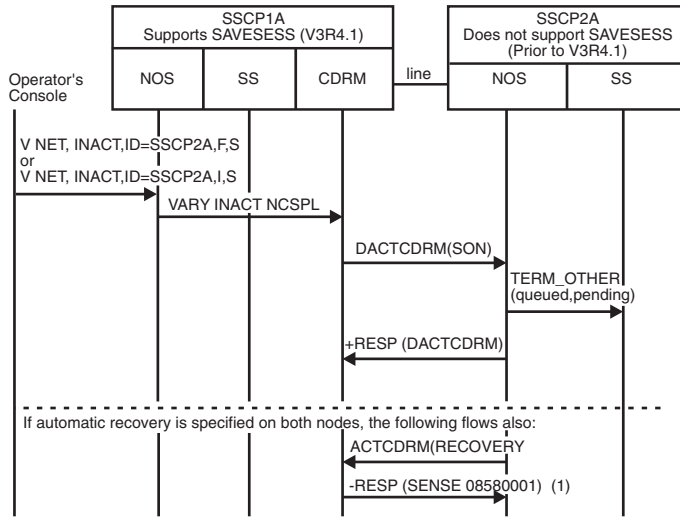


Figure 159. Deactivating a CDRM on a VTAM level before V3R4.1: Forced or immediate

1. When automatic recovery is specified on both nodes, a CDRM V3R4.1 responds to ACTCDRM(RECOVERY) by sending a negative ACTCDRM response with sense code 08580001, indicating that it rejects the attempt to restart the session that was terminated using a nondisruptive deactivation request. Active LU-LU sessions remain active. The external CDRM in the migration SSCP becomes inactive with sessions, and the external CDRM in the V3R4.1 SSCP becomes inactive.

See z/OS Communications Server: IP and SNA Codes for information on sense code 08580001.

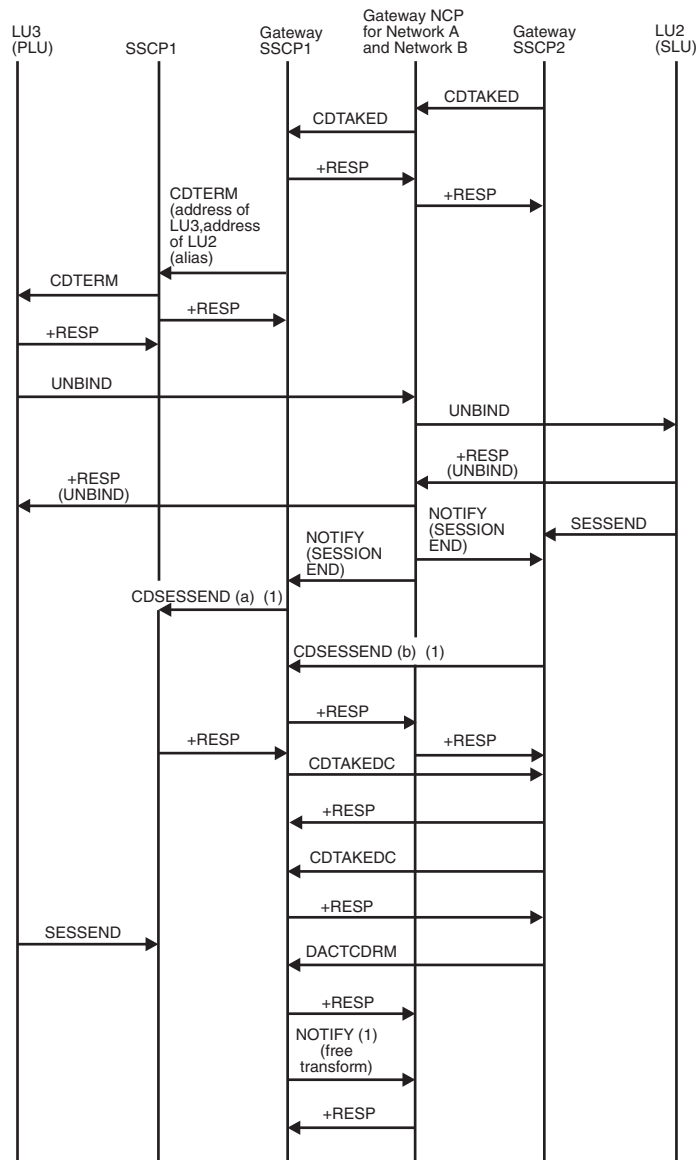


Figure 160. SSCP-SSCP session termination causes LU-LU sessions to be broken

Note:

1. A cross-network LU-LU session exists between LU3 and LU2.
2. This flow assumes that the gateway VTAM1 established the network address translation for the gateway VTAM1-to-gateway VTAM2 session with the RNAA RU. The NOTIFY to free the transform is sent only if the RNAA that established the address specified "retain address."
3. (a) and (b) are used here to differentiate between similar request units.
For details of CDSSESEND processing, see the other flow diagrams listed in "Index of deactivation and session termination flows" on page 492.

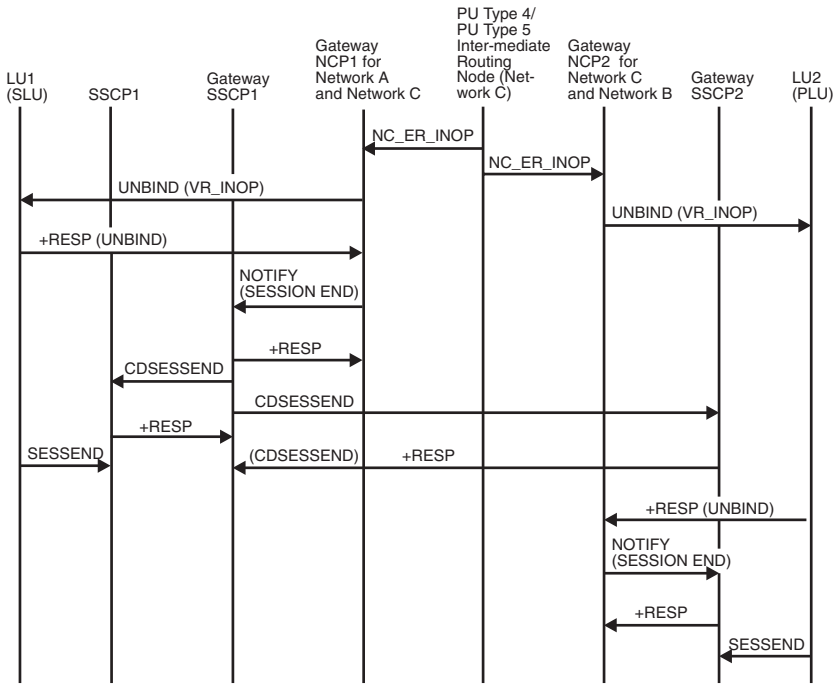


Figure 161. Route failure in intermediate network causes termination of LU-LU sessions

Note: An outage occurs on the route in Network C used by the LU1_LU2 session. ER_INOP reports the failure to gateway NCP1 and gateway NCP2.

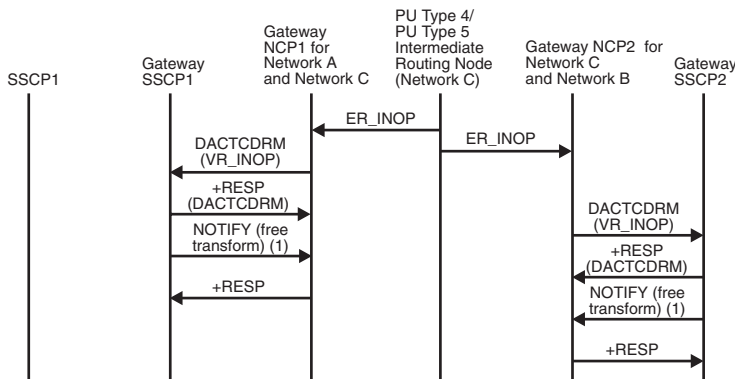


Figure 162. Route failure in intermediate network causes termination of SSCP-SSCP sessions

1. The NOTIFY to free the transform is sent only if the RNAA that established the address specified "retain address."

Note: An outage occurs on the route in Network C used by the gateway VTAM1-to-gateway VTAM2 session. The failure is reported to gateway NCP1 and gateway NCP2 with an ER INOP RU.

Error detection and recovery and SSCP management services

Figure 163 on page 513 through Figure 165 on page 514 show the flow of requests and responses between the SSCP and logical and physical units to handle error recovery processing (ERP) and route Forward and Deliver RUs. Figure 166 on page 514 through Figure 168 on page 516 show the requests and responses between the NetView program, VTAM, the communications adapter, and the local modem for LPDA-2 processing. Figure 169 on page 517 through Figure 172 on page 520 show the flow of requests and responses between the SSCP and logical and physical units to handle extended recovery facility (XRF) session establishments and takeovers with USERVARs.

Index of error detection and recovery and SSCP management services flows

Table 38 lists the error detection and recovery and SSCP management services flows that are illustrated here.

Table 38. Index of error detection and recovery and SSCP management services flows

Flow	Page
Error recovery processing (ERP)	
Hard INOP	Figure 164 on page 513
Soft INOP	Figure 163 on page 513
LPDA-2 processing	
Unsolicited LPDA-2 test on permanent link error with two link segments	Figure 168 on page 516
Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 physical unit (PU) with one link segment	Figure 166 on page 514
Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 physical unit (PU) with two link segments	Figure 167 on page 515
SSCP management services processing	
FORWARD and DELIVER Routing	Figure 165 on page 514
XRF processing	
Secondary logical unit (LU) initiate with USERVAR (LOGON)	Figure 171 on page 519
Third-party initiate (CLSDST PASS)	Figure 172 on page 520
XRF primary and backup sessions, establishment of	Figure 169 on page 517
XRF session switch (takeover)	Figure 170 on page 518

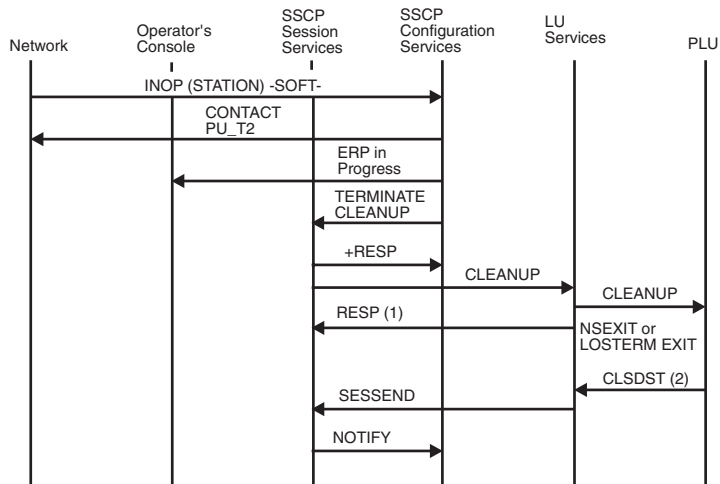


Figure 163. Error recovery processing: Soft INOP

1. If the NSEXIT exit routine is scheduled, LUS cleans up the session and sends a positive response to the cleanup request. If the LOSTERM exit routine is scheduled, LUS *does not* clean up the session, and it sends a negative response to the cleanup request.
2. CLSDST flows only if the LOSTERM exit routine is scheduled.

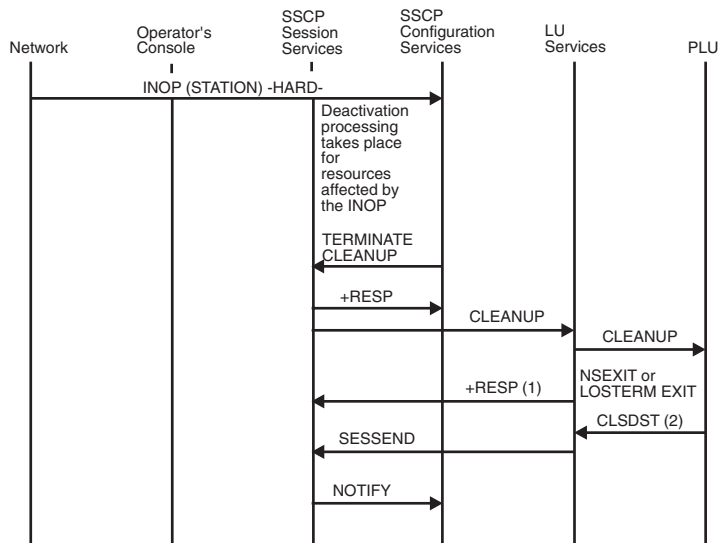


Figure 164. Error recovery processing: Hard INOP

1. If the NSEXIT exit routine is scheduled, LUS cleans up the session and sends a positive response to the cleanup request. If the LOSTERM exit routine is scheduled, LUS *does not* clean up the session, and it sends a negative response to the cleanup request.
2. CLSDST flows only if the LOSTERM exit routine is scheduled.

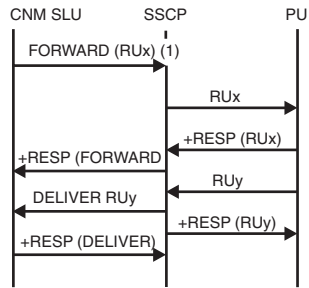


Figure 165. FORWARD and DELIVER routing

1. RUx is a maintenance service RU.
2. Either RUy contains data in reply to RUx, or it is an unsolicited RU.

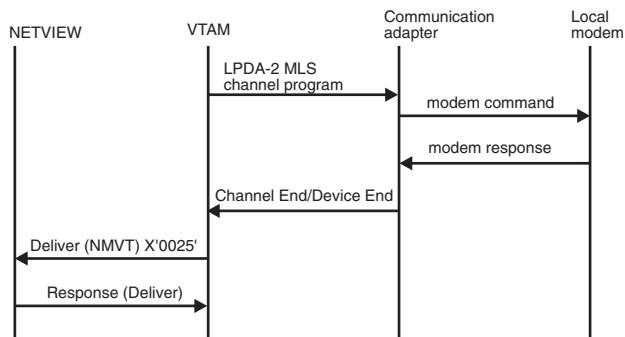


Figure 166. Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 PU with one link segment

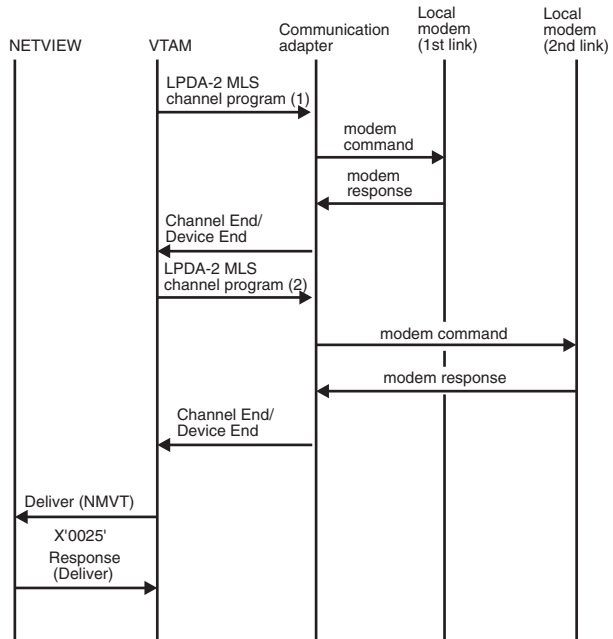


Figure 167. Unsolicited LPDA-2 test on thresholds reached for an LPDA-2 PU with two link segments

1. This MLS (modem and link status) command is for the first link segment.
2. This MLS command is for the second link segment.

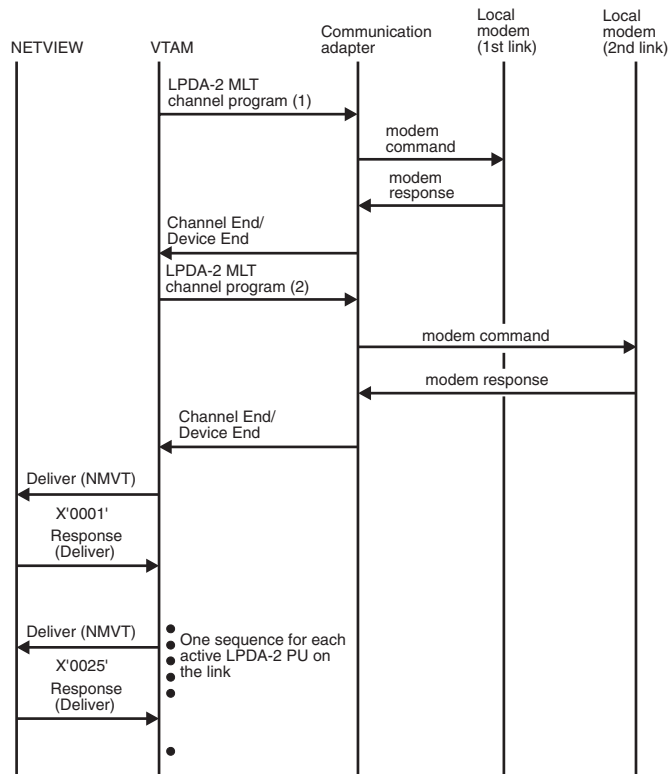


Figure 168. Unsolicited LPDA-2 test on permanent link error with two link segments

1. This MLT (modem and link test) command is for the first link segment.
2. This MLT command is for the first active LPDA-2 PU on the second link segment.

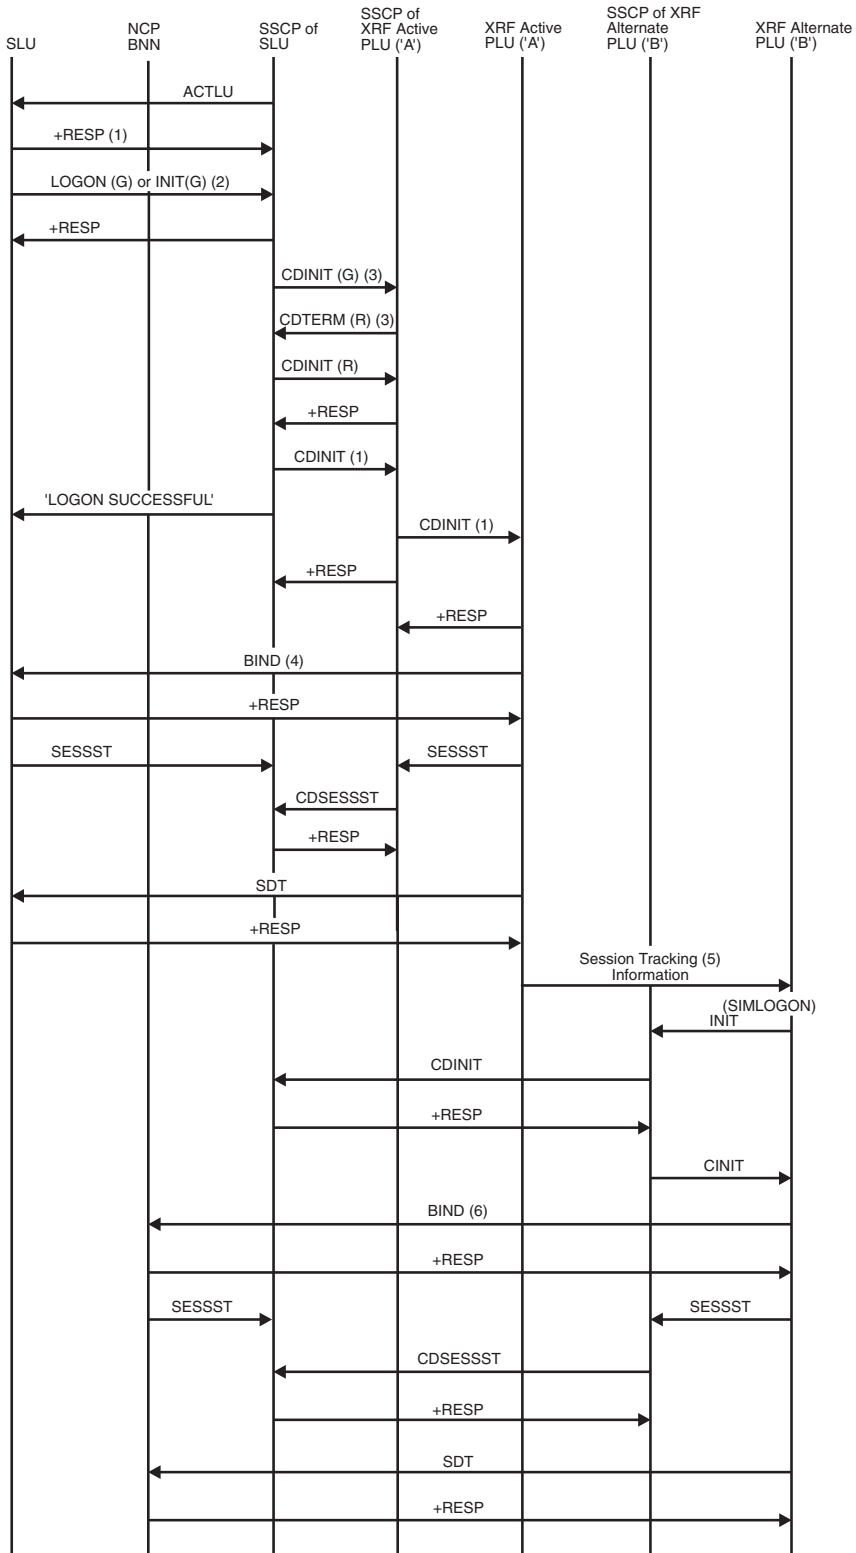


Figure 169. Establishment of XRF primary and backup sessions

- G Represents a generic USERVAR name
- R Represents a resolved USERVAR name

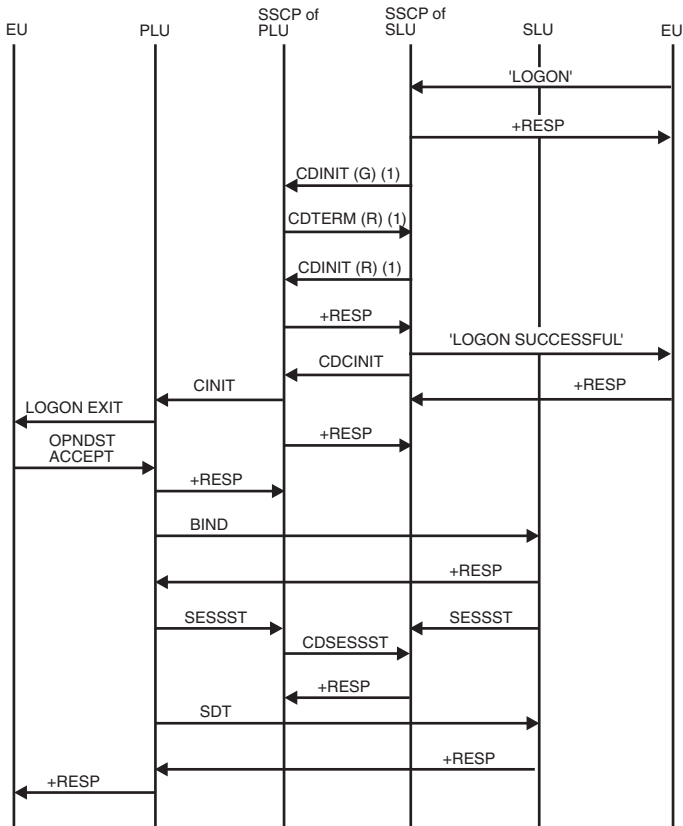


Figure 171. Secondary logical unit initiate with USERVAR (LOGON)

G Represents a generic USERVAR name

R Represents a resolved USERVAR name

1. These RUs are used to translate the generic (USERVAR) name used in the LOGON or INIT_SELF to the real name of the application that is currently the XRF active. They are present only if the SLU's SSCP does not already know the current value of the USERVAR or if the USERVAR's type is VOLATILE.

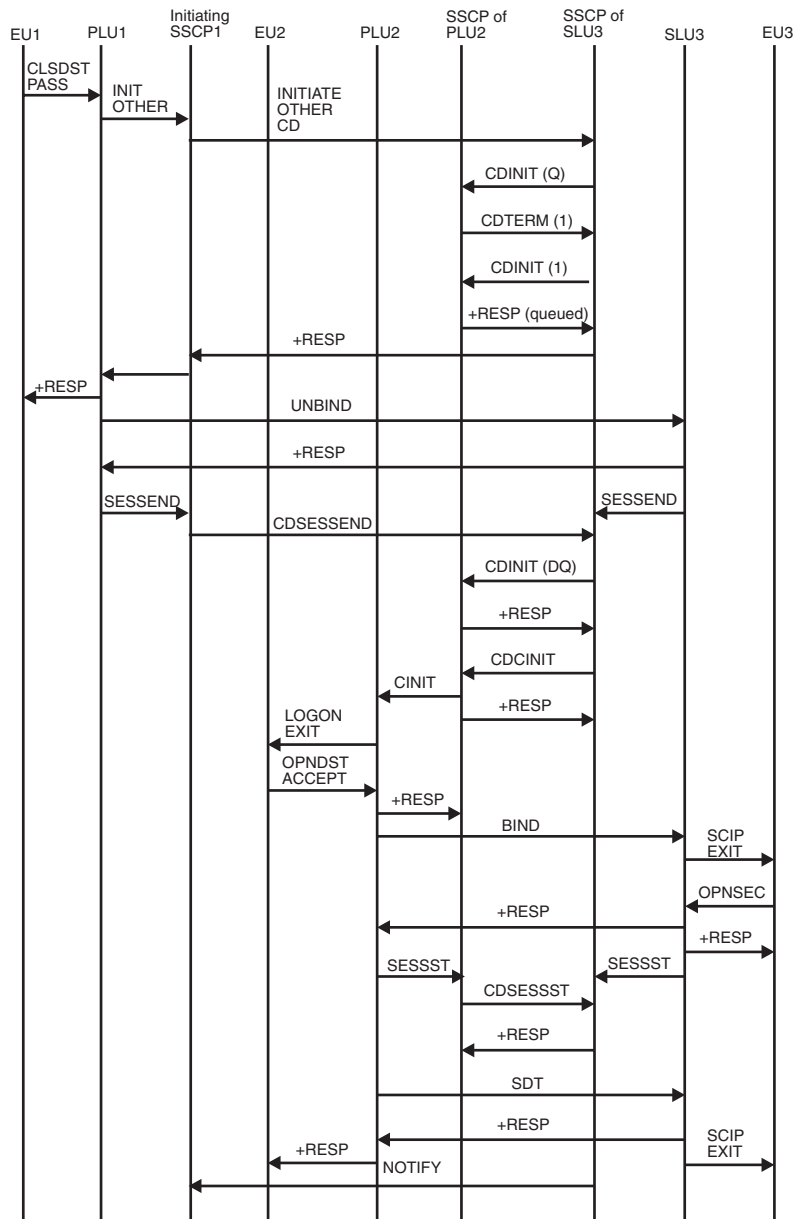


Figure 172. Third-party initiate (CLSDST PASS)

1. These RUs are used to translate the generic (USERVAR) name used in the LOGON or INIT_SELF to the real name of the application. They are present only if the SLU's SSCP does not already know the current value of the USERVAR or if the USERVAR's type is VOLATILE.

Appendix C. APPN flows

This appendix describes the flows between APPN end nodes, network nodes, interchange nodes, and the subarea network. The flow diagrams are divided into the following categories:

- “CP-CP session flows” on page 525
- “Directory services flows” on page 533
- “LU-LU session flows” on page 556
- “Dependent LU server flows” on page 587
- “High-Performance Routing flows” on page 614

Table 39 lists all the APPN flows illustrated here.

Table 39. Index of APPN flows

APPN flow	Page
CP-CP session flow	Page
Activating a CP-CP contention loser session	Figure 174 on page 528
Activating a CP-CP contention winner session	Figure 173 on page 526
Activating an APPN Host-to-Host Channel	Figure 178 on page 532
Activating a leased APPN Node type 2.1	Figure 177 on page 531
Host CP initiating deactivation of CP-CP session	Figure 175 on page 529
Remote node initiating deactivation of CP-CP session	Figure 176 on page 530
<hr/>	
Dependent LU server flow	Page
Single subnetwork	
Activating resources	
CPSVRMGR pipe activation, DLUR-Initiated	Figure 238 on page 589
CPSVRMGR pipe activation, DLUS-initiated	Figure 239 on page 590
Dependent LUs, dynamic registration and activation of	Figure 241 on page 592
Dependent LUs, activation of pre-defined	Figure 242 on page 592
	Figure 240 on page 591
SSCP-PU session activation race	Figure 243 on page 593
Deactivating resources	
CPSVRMGR pipe deactivation	Figure 244 on page 594
Downstream PU outage	Figure 245 on page 595
REQDISCONT (immediate) received from downstream PU	Figure 247 on page 597
REQDISCONT (normal) received from downstream PU	Figure 246 on page 596
LU-LU sessions	
APPN PLU-initiated to a dependent SLU	Figure 252 on page 602
Session termination, USS flows for	Figure 255 on page 605
USS SLU-initiated to APPN PLU	Figure 253 on page 603
USS SLU-initiated to subarea PLU	Figure 254 on page 604

Table 39. Index of APPN flows (continued)

APPN flow	Page
SSCP-PU, SSCP-LU session deactivation	
Forced	Figure 249 on page 599
Normal	Figure 248 on page 598
With Giveback (ANS=CONT)	Figure 251 on page 601
With Giveback (ANS=STOP)	Figure 250 on page 600
Cross subnetwork	
PLU-Initiated Session with DLUS and DLUR within Different Subnetworks	Figure 256 on page 606
PLU-Initiated Session with DLUS and PLU in one Subnetwork and DLUR in Another	Figure 258 on page 609
SLU-Initiated Session with DLUS and DLUR within Different Subnetworks	Figure 260 on page 611
Directory services flow	
Locate resource	
APPN and subarea Network	Figure 192 on page 548
APPN network, complex	Figure 187 on page 541
Complex APPN network using more than one CDS	Figure 188 on page 543
CP-CP session terminates	Figure 197 on page 554
CP network broadcast initiation	Figure 185 on page 539
EN to NN	Figure 181 on page 536
EN to NN to EN	Figure 182 on page 536
EN to NN to NN to NN	Figure 184 on page 538
EN to NN to subarea network	Figure 186 on page 540
EN to NN to two ENs	Figure 183 on page 537
NN receives network broadcast request	Figure 198 on page 555
NNS of the OLU is at pre-V4R2 level	Figure 190 on page 546
Directory search verification reduction	Figure 195 on page 551
SLU-initiated session	Figure 196 on page 552
Register resource	
EN to NN to CDS	Figure 179 on page 534
With error recovery	Figure 180 on page 535
High Performance Routing (HPR) flow	
Rapid-transport protocol (RTP) connection over portion of session path	Figure 264 on page 617
Rapid-transport protocol (RTP) across composite nodes with a T2.1 connection through NCP	Figure 265 on page 618
Rapid-transport protocol (RTP) across composite nodes with a T2.1 connection through VTAM	Figure 266 on page 619
Rapid-transport protocol (RTP) across composite nodes with a virtual-route-based transmission group, NCP does ANR routing	Figure 267 on page 620

Table 39. Index of APPN flows (continued)

APPN flow	Page
Rapid-transport protocol (RTP) across composite nodes with a virtual-route-based transmission group, VTAM does ANR routing	Figure 268 on page 621
Two rapid-transport protocol (RTP) nodes with a T2.1 connection	Figure 262 on page 615
Two rapid-transport protocol (RTP) nodes with a virtual-route-based transmission group	Figure 263 on page 616
LU-LU sessions flow	Page
APPN network...NNS--EN (PLU)	
SLU-initiated, no queueing	Figure 206 on page 562
SLU-initiated, queued by the PLU	Figure 207 on page 562
APPN network...NNS--EN(SLU)	
PLU-initiated, no queueing	Figure 204 on page 561
PLU-initiated, queued by the SLU	Figure 205 on page 561
APPN network (PLU)...ICN==SA(SLU), PLU-Initiated	
Directed search without required precomputed RSCV	Figure 230 on page 580
No queueing	Figure 229 on page 579
Queued by SLU	Figure 232 on page 582
Search-only flow transformed into a DSRLST	Figure 228 on page 578
USERVAR resolution required	Figure 231 on page 581
APPN network (PLU)...ICN==VR-based TG==ICN...APPN network (SLU)	
PLU-initiated	Figure 237 on page 587
APPN network (SLU)...ICN==SA(PLU)	
Autologon, PLU not available initially	Figure 236 on page 586
SLU-initiated, no queueing	Figure 233 on page 583
SLU-initiated, queued by the PLU	Figure 234 on page 584
CLSDST PASS; SLU is single-session capable	
From APPN to subarea	Figure 225 on page 576
Through APPN	Figure 224 on page 575
EN-NN-EN, PLU-initiated, no queueing (including BIND flows for intermediate network node)	Figure 226 on page 577
EN (PLU)--NNS...APPN network	
PLU-initiated, no queueing	Figure 199 on page 558
PLU-initiated, queued by the PLU	Figure 200 on page 559
PLU-initiated, queued by the SLU	Figure 201 on page 559
EN (SLU)--NNS...APPN network	
SLU-initiated, no queueing	Figure 202 on page 560
SLU-initiated, queued by the PLU	Figure 203 on page 560
Intermediate Network Node (INN) BIND. The LOCATE did not go through this node.	Figure 227 on page 578
SA (PLU)==ICN...APPN network (SLU)	

Table 39. Index of APPN flows (continued)

APPN flow	Page
DSRLIST transforming into PLU-initiated, search-only	Figure 208 on page 563
PLU-initiated, no queueing	Figure 209 on page 564
PLU-initiated, queued by the SLU	Figure 211 on page 565
PLU-initiated, USERVAR resolution required	Figure 210 on page 565
SA (SLU)==ICN...APPN network (PLU)	
Autologon, PLU not available initially	Figure 215 on page 569
SLU-initiated, no queueing	Figure 213 on page 567
SLU-initiated, queued by the PLU	Figure 214 on page 568
Session release request	
SA(PLU)==ICN...APPN network(SLU)	Figure 222 on page 574
SA(SLU)==ICN...APPN network(PLU)	Figure 223 on page 574
Session termination, forced	
SA(PLU)==ICN...APPN Network(SLU), pending active session. PLU is accessible without going into APPN.	Figure 219 on page 572
SA(PLU)==ICN...APPN Network(SLU), queued session	Figure 220 on page 573
SA(PLU)==ICN...APPN Network(SLU), queued session. PLU is accessible without going into APPN.	Figure 221 on page 573
SA(SLU)==ICN...APPN Network(PLU), pending active session	Figure 218 on page 571
Session termination, orderly	
SA(PLU)==ICN...APPN Network(SLU), active session	Figure 216 on page 570
APPN Network(PLU)...ICN==SA(SLU), active session	Figure 217 on page 571

Many abbreviations are shown at the top of the flow diagrams. The following list gives the meaning of those abbreviations:

- ANR** Automatic network routing
- APPC** Advanced program-to-program communication program
- BF** Boundary function
- CDS** Central Directory Server
- CP** Control point
- DLUR** Dependent logical unit requestor
- DLUS** Dependent logical unit server
- EN** End node
- ICN** Interchange node
- LU** Logical unit
- NN** Network node
- NNS** Network node server
- PLU** Primary logical unit
- PU** Physical unit

PUNS	Physical unit services
RTP	Rapid-transport protocol
SLU	Secondary logical unit
SSCP	System services control point
TP	Transaction program
TSC	Transmission subsystem component

CP-CP session flows

This information illustrates communication protocols between nodes in a mixed APPN and subarea network. Use these flows as guidelines to help analyze and isolate network problems caused by unexpected network events, such as protocol violations.

Index of CP-CP session flows

Table 40 lists the CP-CP session flows illustrated here.

Table 40. Index of CP-CP session flows

Flow	Page
Activating a CP-CP contention loser session	Figure 174 on page 528
Activating a CP-CP contention winner session	Figure 173 on page 526
Activating an APPN Host-to-Host Channel	Figure 178 on page 532
Activating a leased APPN Node type 2.1	Figure 177 on page 531
Host CP initiating deactivation of CP-CP session	Figure 175 on page 529
Remote node initiating deactivation of CP-CP session	Figure 176 on page 530

Activating a CP-CP contention-winner session

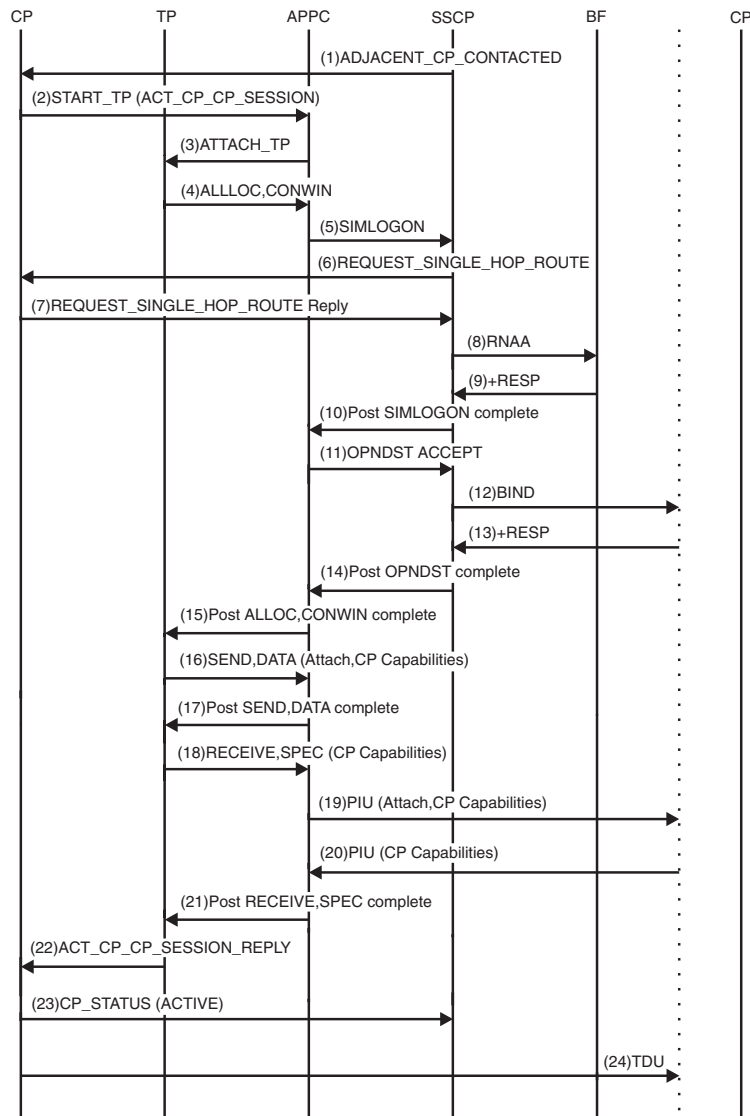


Figure 173. CP-CP contention-winner session activation

1. Configuration services sends an ADJACENT_CP_CONTACTED signal to the CP when the first link supporting CP-CP sessions is activated with an adjacent CP.
2. The CP sends a START_TP signal to the APPC PAB to request initiation of the Request CP Capabilities Transaction Program (TP). An ACT_CP_CP_SESSION request is queued to the START_TP signal for processing by the Request CP Capabilities TP. The ACT_CP_CP_SESSION initiates the activation of a contention winner CP-CP session with an adjacent node.
3. APPC sends an ATTACH_TP AMU for the Request CP Capabilities TP to TP Services.
4. The Request CP Capabilities TP issues an APPCCMD CONTROL=ALLOC, QUALIFY=CONWIN macroinstruction to allocate a conversation between the TP and the partner TP and a contention winner session between the local LU and the remote LU.

5. APPC issues a VTAM SIMLOGON macroinstruction to initiate a session in which APPC acts as the PLU.
6. Subarea session services sends a REQUEST_SINGLE_HOP_ROUTE signal to the CP to request the least-weight single-hop route from the origin to the destination.
7. The CP sends to subarea session services the information requested in a REQUEST_SINGLE_HOP_ROUTE_REPLY.
8. RNAA flows if the CP-CP session is being set up over a type 2.1 link and a network address is needed. If the CP-CP session is being set up over a VR-based transmission group, CDINIT format 5 is sent to the session partner to get a network address.
9. The response is received from the boundary function.
10. LUS posts the SIMLOGON complete.
11. APPC issues an OPNDST ACCEPT macroinstruction to continue establishment of a session between APPC in this node (acting as the PLU) and APPC in the adjacent node (acting as the SLU).
12. The BIND for the contention winner session is transmitted to the adjacent node.
13. The BIND response for the contention winner session is received from the adjacent node.
14. LUS posts the OPNDST ACCEPT complete.
15. APPC posts the APPCCMD CONTROL=ALLOC,QUALIFY=CONWIN complete, supplying the Request CP Capabilities TP a conversation ID and a contention winner conversation group ID (CGID).
16. The Request CP Capabilities TP issues an APPCCMD CONTROL=SEND, QUALIFY=DATA macroinstruction to initiate the sending of the CP Capabilities data to the adjacent CP.
17. APPC posts the APPCCMD CONTROL=SEND,QUALIFY=DATA instruction complete, indicating that the output buffer has been filled with the CP Capabilities data.
18. At the request of the Request CP Capabilities TP, the Receive and Check CP Capabilities TP issues an APPCCMD CONTROL=RECEIVE, QUALIFY=SPEC macroinstruction to cause the transmission of the CP Capabilities data to the adjacent CP, and to initiate the receiving of CP Capabilities data from the adjacent CP.
19. APPC sends to the adjacent node a PIU with the CP Capabilities data. The PIU also carries a request that TP services in the adjacent node attach its CP Capabilities TP.
20. APPC receives from the adjacent node a PIU containing the CP Capabilities of the adjacent CP.
21. APPC posts the APPCCMD CONTROL=RECEIVE,QUALIFY=SPEC macroinstruction complete.
22. The Request CP Capabilities TP responds to the successful completion of the RECEIVE macroinstruction by sending to the CP an ACT_CP_CP_SESSION_REPLY, which contains both the contention winner CGID and the CP Capabilities data received from the adjacent CP.
23. If both the contention winner and contention loser CP-CP sessions are active, the CP sends a CP_STATUS(ACTIVE,BOTH) signal to configuration services.

24. If both CPs are network nodes, a topology database update (TDU) will flow when the contention winner session is active. The TDU is used to update the partner regarding changes to network topology that have occurred since the two CPs were last in session.

If the network node server is a VTAM, a TDU will also flow over the contention winner session from a VTAM end node to its network node server. This TDU carries information about changes that have occurred to the end node connections since the end node and the network node server were last in session.

Activating a CP-CP contention-loser session

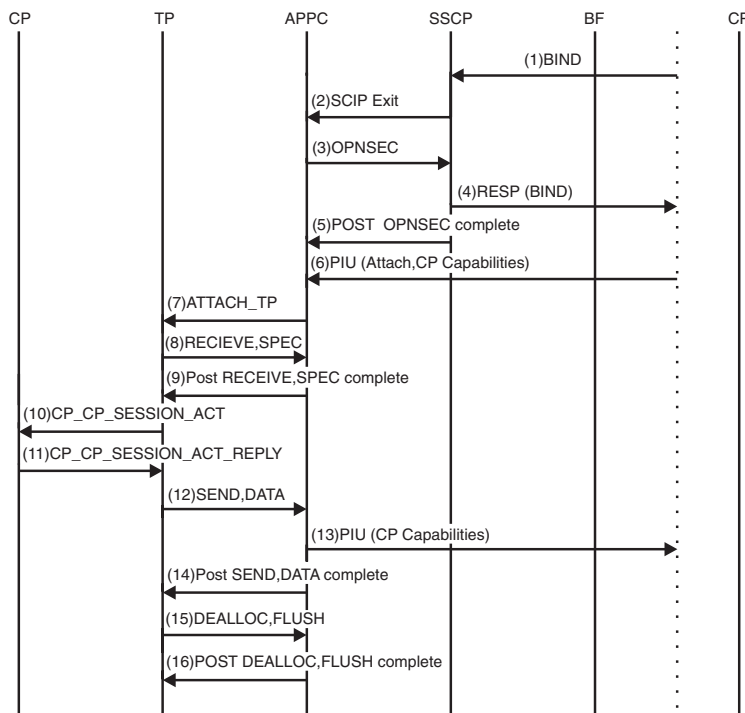


Figure 174. CP-CP contention-loser session activation

1. The BIND for the contention loser CP-CP session is received from the adjacent node.
2. The BIND drives the SCIP exit routine.
3. The SCIP exit causes APPC to issue an OPNSEC macroinstruction to establish a session between APPC (operating as the SLU) and the PLU that sent the BIND.
4. The OPNSEC macroinstruction causes a BIND response to be sent to the adjacent node.
5. The OPNSEC is posted complete.
6. APPC receives from the adjacent node a PIU containing the CP Capabilities of the adjacent CP. The PIU also contains an FMH-5 specifying that the CP Capabilities Transaction Program (TP) is to be attached.
7. APPC sends an ATTACH_TP AMU to TP Services to request the attachment of the CP Capabilities TP.

8. After attachment, the CP Capabilities TP issues an APPCCMD CONTROL=RECEIVE,QUALIFY=SPEC instruction to receive the CP Capabilities of the adjacent CP.
9. APPC posts the APPCCMD CONTROL=RECEIVE,QUALIFY=SPEC macroinstruction complete.
10. The CP Capabilities TP builds and sends to the CP a CP_CP_SESSION_ACT signal, containing the contention loser Conversation Group Identifier (CGID) and the CP Capabilities of the adjacent CP.
11. The CP sends to the CP Capabilities TP a CP_CP_SESSION_ACT_REPLY. It contains the CP Capabilities of this node and also indicates whether SSC was able to successfully process the CONLOSER activation request.
12. After receiving the CP_CP_SESSION_ACT_REPLY, the CP Capabilities TP issues an APPCCMD CONTROL=SEND,QUALIFY=DATA to send the CP Capabilities to the adjacent CP.
13. A PIU with the CP Capabilities is transmitted to the adjacent CP.
14. When the transmission is complete, the APPCCMD CONTROL=SEND,QUALIFY=DATA macroinstruction is posted complete by APPC.
15. The CP Capabilities TP issues an APPCCMD CONTROL=DEALLOC,QUALIFY=FLUSH macroinstruction to flush the local LU send buffer and deallocate the conversation normally.
16. APPC posts the APPCCMD CONTROL=DEALLOC,QUALIFY=FLUSH macroinstruction back upon completion.

Host CP initiating deactivation of CP-CP session

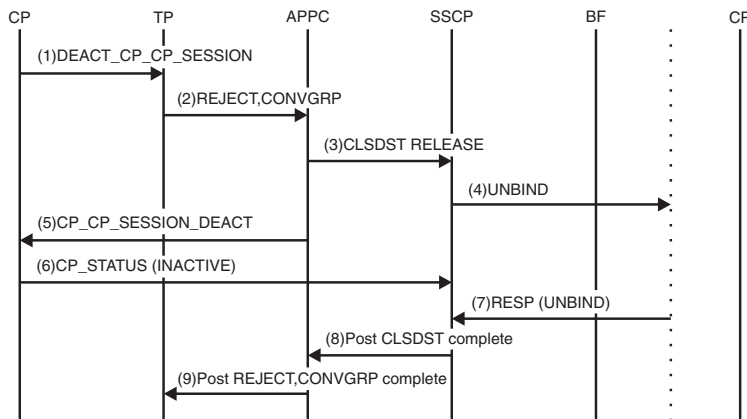


Figure 175. Host CP initiating deactivation of CP-CP session

1. The CP sends a DEACT_CP_CP_SESSION to the SEND_REJECT_CONVGRP Transaction Program (TP). The CP sends this signal when it must have a particular CP-CP session with an adjacent node unbound.
2. The SEND_REJECT_CONVGRP TP issues an APPCCMD CONTROL=REJECT,QUALIFY=CONVGRP macroinstruction to deactivate the specified session.
3. APPC, responding to the APPCCMD CONTROL=REJECT,QUALIFY=CONVGRP macroinstruction, issues a CLSDST RELEASE macroinstruction to terminate the session.

4. APPC also builds and sends a CP_CP_SESSION_DEACT for the specified session to the CP. This signal is sent by APPC to notify the CP that a CP-CP session outage is detected. It contains the session type and CGID of the CP-CP session to which the outage is detected.
5. The CLSDST RELEASE causes SSCP to send an UNBIND for the particular session to the partner LU.
6. The CP sends a CP_STATUS(INACTIVE) signal for the specified session to the SSCP.
7. A response to the UNBIND is received by the SSCP from the partner LU.
8. Having received the UNBIND response, the SSCP posts complete the CLSDST RELEASE.
9. APPC posts the TP APPCCMD CONTROL=REJECT,QUALIFY=CONVGRP macroinstruction complete.

Remote node initiating deactivation of CP-CP session

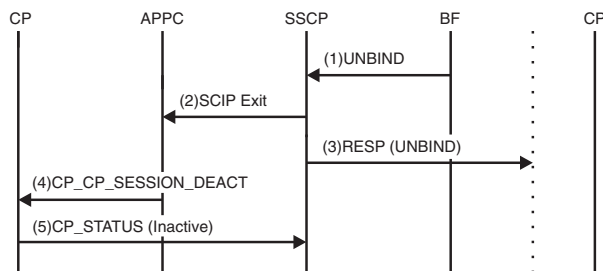


Figure 176. Remote node initiating deactivation of CP-CP session

1. An UNBIND is received from the adjacent node for a specific CP-CP session. The UNBIND carries a sense code associated with the session outage.
2. The UNBIND drives the SCIP exit routine.
3. The SSCP sends an UNBIND response to the adjacent node.
4. APPC sends a CP_CP_SESSION_DEACT to the CP for the session specified in the UNBIND. The CP_CP_SESSION_DEACT carries the sense code originally carried by the UNBIND.
5. The CP notifies the SSCP that the specified session is now inactive.

Activating a leased APPN node type 2.1

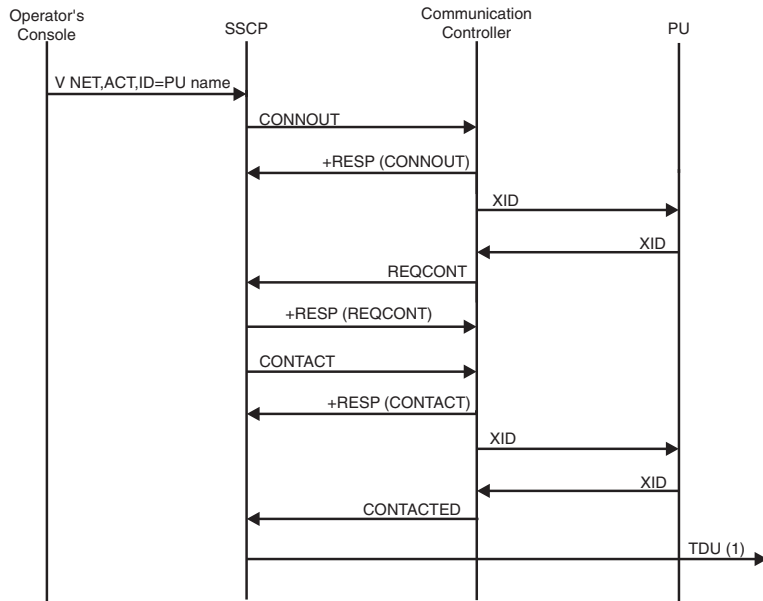


Figure 177. Activating a leased APPN node type 2.1

If this is a network node, a topology database update (TDU) is sent to all CP-CP session partners, informing them that this new APPN connection is available. If this is an end node, and the network node server is a VTAM, a TDU will be sent to the server, informing it of the new connection.

Leased APPN PUs, as opposed to non-APPN PUs, have an additional CONNOUT, XID, and REQCONT flow. This flow allows transmission group (TG) negotiation during prenegotiation XID exchanges.

Activating an APPN host-to-host channel

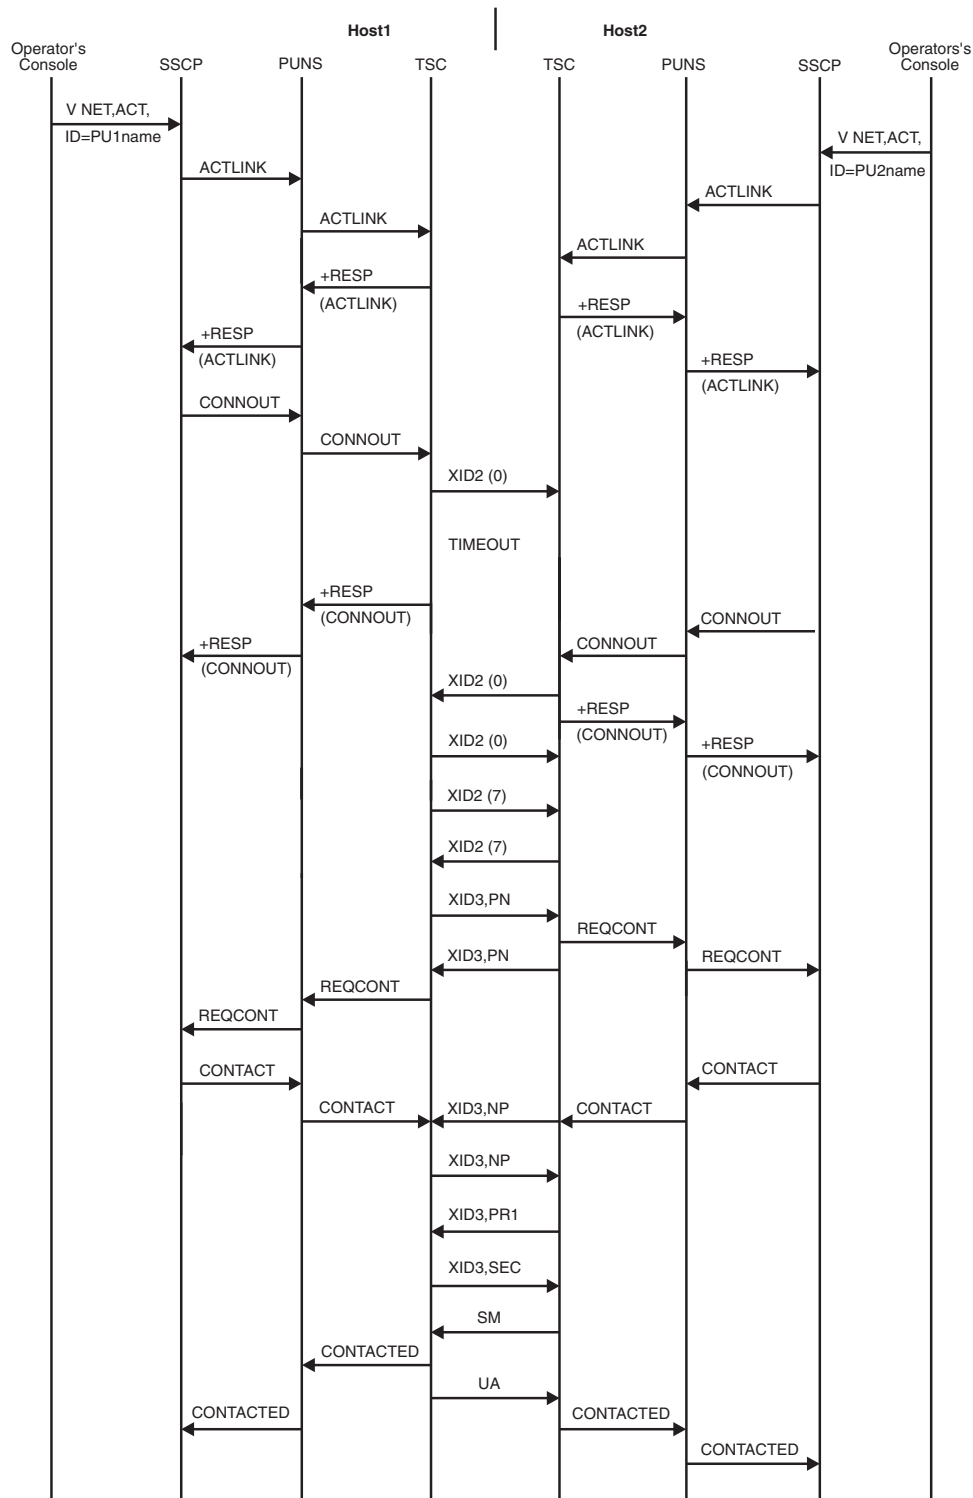


Figure 178. Activating an APPN host-to-host channel

Note: XID2 Type 0 and Type 7 are sent across each subchannel associated with an APPN host-to-host channel connection. The XID3s are sent by choosing one of the available write subchannels for transmission.

Directory services flows

To understand the following flows, it is helpful to understand the concept of resource registration and the implications of the values coded on the REGISTER operand. For information on resource registration and registering application programs, see *z/OS Communications Server: SNA Network Implementation Guide*. For information about coding the REGISTER operand, see *z/OS Communications Server: SNA Resource Definition Reference*.

In these directory services flows, assume the following conditions, unless stated otherwise:

- Sessions are initiated by the primary logical unit (PLU).
- The flows illustrate search-only requests.

Index of directory services flows

Table 41 lists the directory services flows illustrated here.

Table 41. Index of directory services flows

Flow	Page
Locate resource	
APPN and subarea Network	Figure 192 on page 548
APPN network, complex	Figure 187 on page 541
Complex APPN network using more than one CDS	Figure 188 on page 543
CP-CP session terminates	Figure 197 on page 554
CP network broadcast initiation	Figure 185 on page 539
EN to NN	Figure 181 on page 536
EN to NN to EN	Figure 182 on page 536
EN to NN to NN to NN	Figure 184 on page 538
EN to NN to subarea network	Figure 186 on page 540
EN to NN to two ENs	Figure 183 on page 537
NN receives network broadcast request	Figure 198 on page 555
NNS of the OLU is at pre-V4R2 level	Figure 190 on page 546
Directory search verification reduction	Figure 195 on page 551
SLU-initiated session	Figure 196 on page 552
Resource registration flows	
EN to NN to CDS	Figure 179 on page 534
With error recovery	Figure 180 on page 535

Register resource flows

Figure 179 on page 534 and Figure 180 on page 535 show the process of registering resources.

Resource registration: EN to NN to CDS

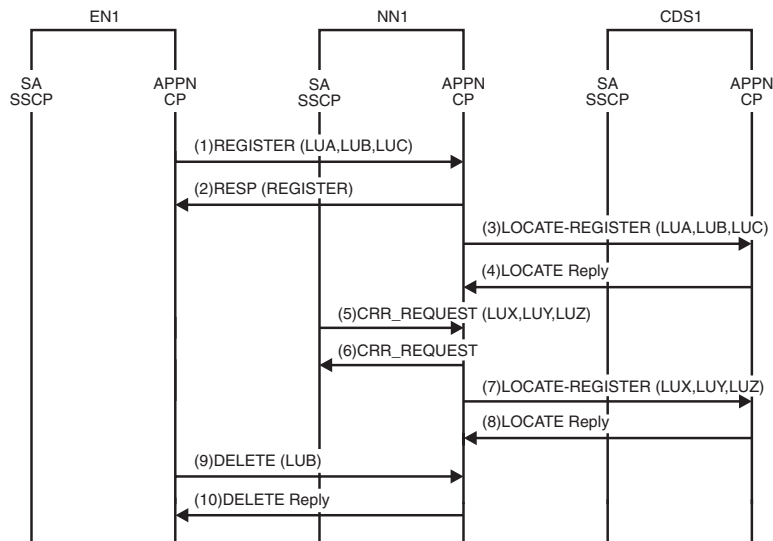


Figure 179. Resource registration: EN to NN to CDS

Note: In this figure, the end node registers its resources to NN1:

- LUA
- LUB
- LUC

NN1 owns the following resources:

- LUX
- LUY
- LUZ

NN1 registers the following resources to the central directory server (CDS):

- LUA
- LUB
- LUC
- LUX
- LUY
- LUZ

1. The VTAM operator in the end node activates a major node containing LUA, LUB, and LUC. The end node sends a registration request to NN1.
2. NN1 adds entries to the directory database and then sends a registration reply to EN1.
3. NN1 sends a central registration request to CDS1. The registration request travels with a LOCATE GDS variable because the central directory server can be several nodes away.
4. CDS1 returns a reply to NN1.
5. The VTAM operator activates a major node containing applications X, Y, and Z. The subarea SSCP notifies the APPN control point (CP) of resources owned by NN1 that should be centrally registered. (CRR stands for central resource registration.)
6. APPN CP sends an immediate reply to the subarea SSCP.

7. NN1 sends a central registration request to CDS1.
8. CDS1 returns a reply to NN1.
9. The end node sends a DELETE request to NN1.
10. NN1 removes LUB from its directory database and returns a reply to the end node. NN1 does not forward the DELETE request to the central directory server.

Resource registration with error recovery

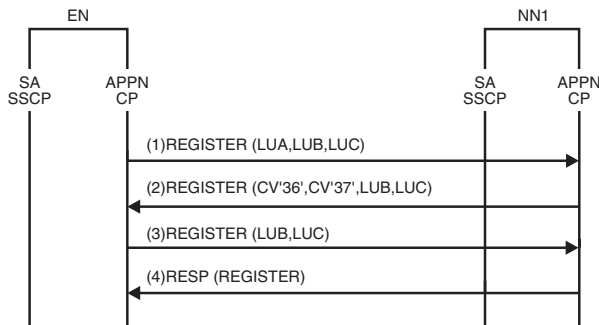


Figure 180. Resource registration with error recovery

1. The VTAM operator in the end node activates a major node containing LUA, LUB, and LUC. The end node sends a registration request to NN1.
2. The network node begins adding resources to its directory database. The network node successfully adds LUA. However, it encounters a problem and cannot continue adding resources to the directory database. The network node sends a negative reply to the end node to indicate which resource the network node was trying to add when it encountered the problem. The CV'36' indicates the sense code. The CV'37' indicates where the network node stopped adding to its directory database.
3. The end node tries again to register those resources that were not successfully registered before.
4. The network node successfully adds to the directory database and returns a reply.

Locate resource flows

Figure 181 on page 536 through Figure 198 on page 555 show the search process.

Locate resource: EN to NN

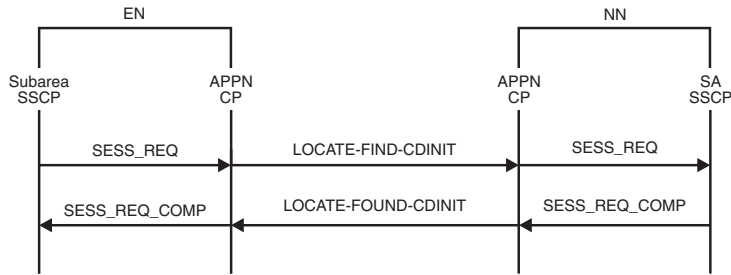


Figure 181. Locate resource: EN to NN

Locate resource: EN to NN to EN

CONFIG: EN — NN — EN

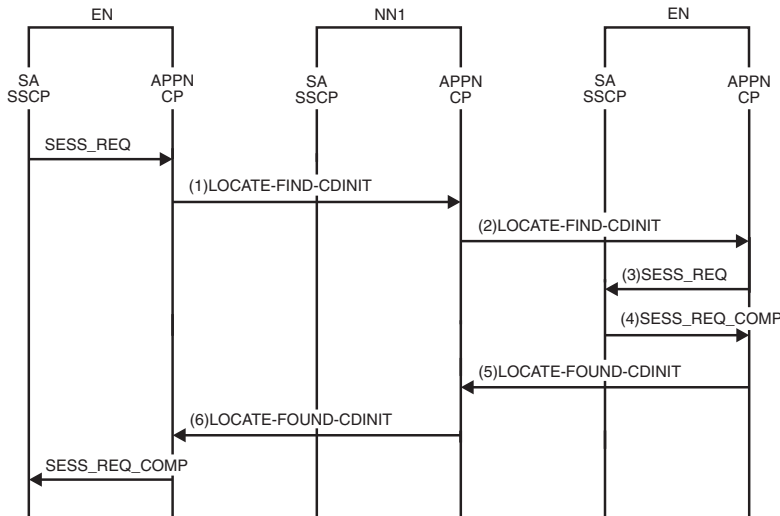


Figure 182. Locate resource: EN to NN to EN

Note: Nodes are connected by CP-CP sessions only.

1. The end node sends a search request for a target resource to the network node server. As the network node server of the originating LU, NN1 looks for the target resource in the directory database. NN1 has knowledge in the directory that the target resource resides on a served end node.
2. NN1 sends a search request to the end node.
3. The APPN control point (CP) sends a SESS_REQ signal to the SSCP.
4. The SSCP sends a SESS_REQ_COMP signal to the APPN CP, indicating that the target resource is located.
5. The end node sends a LOCATE reply to the network node server.
6. NN1 sends a LOCATE reply to the end node.

Locate resource: EN to NN to two ENs

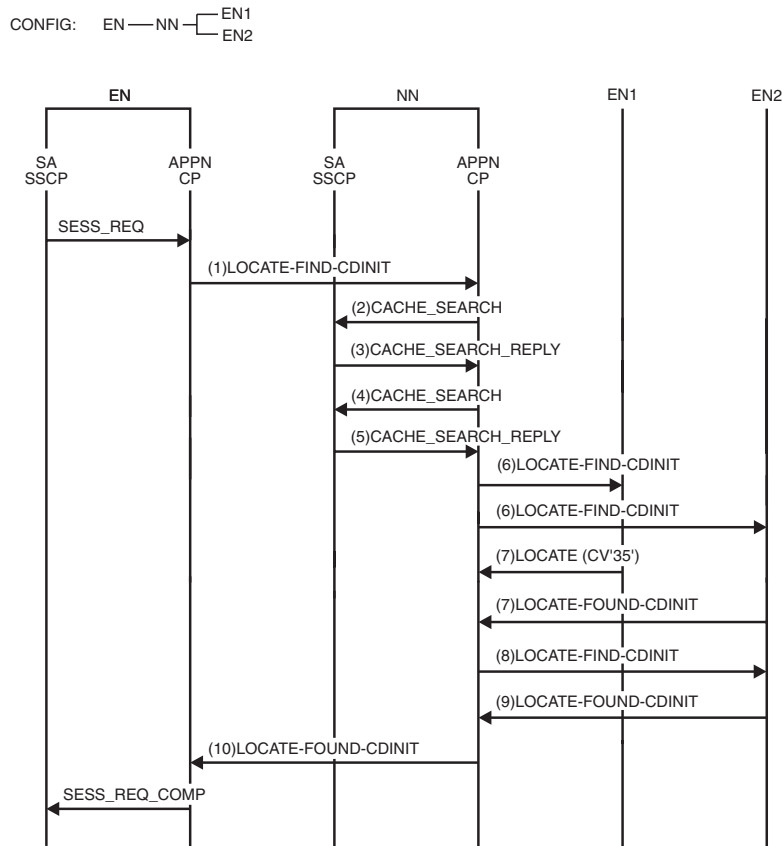


Figure 183. Locate resource: EN to NN to two ENs

Note: Nodes are connected by CP-CP sessions only.

1. The end node sends a search request for a target resource to the network node server. The network node server looks for the target resource in the directory database. The end node that owns the resource did not register its resources.
2. The APPN control point (CP) requests that the subarea SSCP check for information about the location of the resource.
3. The subarea SSCP replies that the target resource is not known.
4. Because NN does not have the target resource in either its APPN or subarea directories, it initiates a resource discovery search. The resource discovery search starts at the beginning of the search logic with a generic request, which is not linked to the original OLU. Because the resource discovery search starts at the beginning of the search logic, another CACHE_SEARCH is performed.
5. The subarea SSCP replies that the target resource is not known.
6. The network node server performs a domain broadcast by sending the search request to all served end nodes that indicate on the CP_CAPABILITIES exchange that they are to be searched on domain broadcast.
7. Each end node that receives the request replies to the search request. EN1 replies that the resource is not found. EN2 replies that it owns the resource.
8. Because the resource discovery search found the resource, the NN sends a search to the target, containing the original session-specific information.
9. The target is found.
10. The network node server replies to the end node.

Locate resource: EN to NN to NN to NN

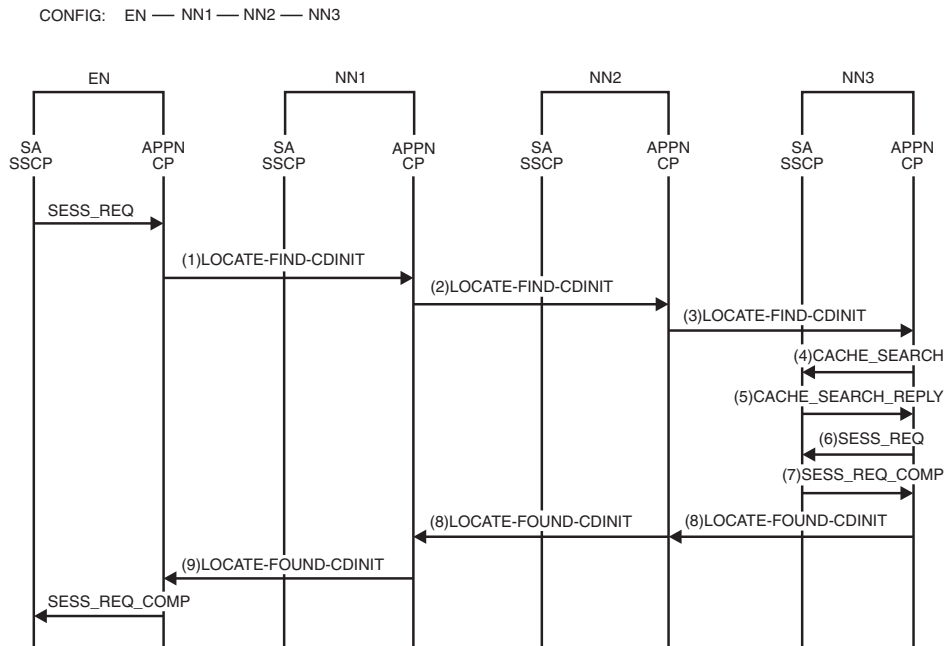


Figure 184. Locate resource: EN to NN to NN to NN

Note: Nodes are connected by CP-CP sessions only.

1. The end node sends a search request for a target resource to the network node server. The network node looks for the target resource in the directory database. The network node has knowledge in the directory that the target resource resides on NN3.
2. The network node sends a directed search request to NN3. Because NN1 does not have direct CP-CP sessions with NN3, NN1 sends the directed search request to NN3 through NN2.
3. NN2 is not the destination of the directed search; therefore, NN2 forwards the request to NN3.
4. APPN control point (CP) sends a request for information to the subarea SSCP.
5. The subarea SSCP replies that the target resource can be found in the subarea network.
6. APPN control point (CP) sends a SESS_REQ signal to the subarea SSCP.
7. The subarea SSCP sends a SESS_REQ_COMP signal to the APPN CP, indicating that the target resource is located.
8. NN3 sends a LOCATE reply to NN2, which forwards the reply to NN1.
9. NN1 sends a LOCATE reply to the end node.

Locate resource: CP network broadcast initiation

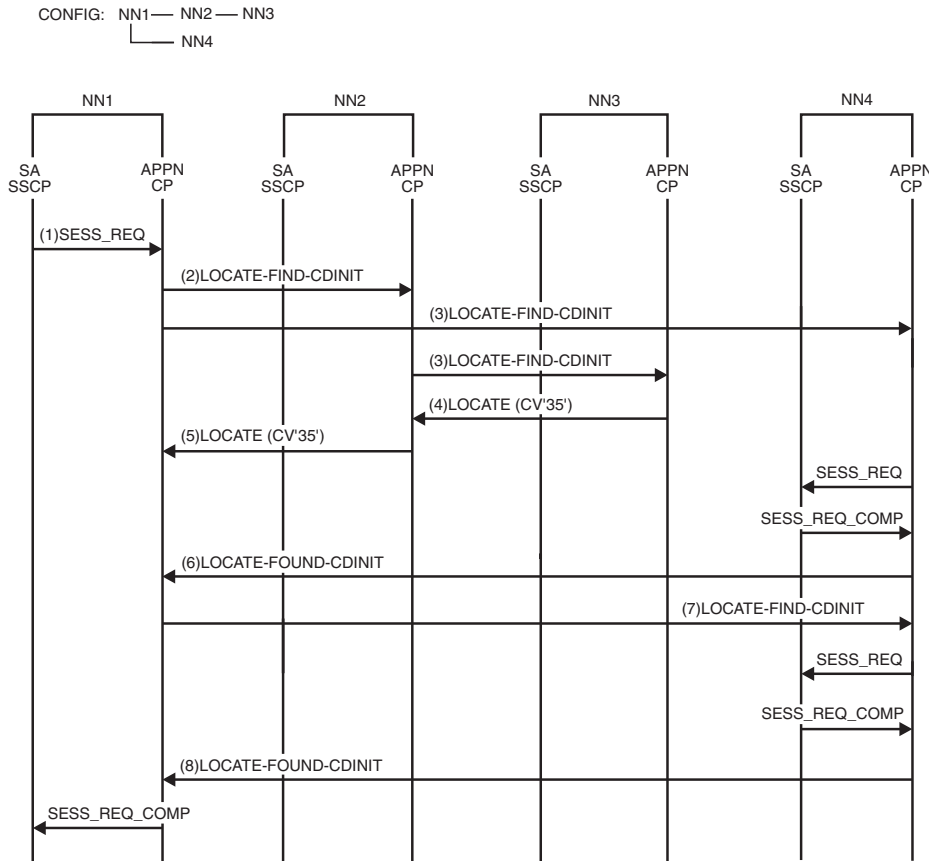


Figure 185. Locate resource: CP network broadcast initiation

1. A resource in NN1 requests a search for a resource on NN4. As network node server of the originating LU, NN1 looks for the target resource in the directory database. NN1 has no knowledge of the location of the target resource. NN1 initiates a resource discovery search for the target, which contains no session-specific information.
2. NN1 has no APPN-domain end nodes, therefore no domain broadcast occurs. There is no central directory server in the network; therefore, NN1 sends a broadcast search request to every network node with which NN1 has CP-CP sessions.
3. Each network node that receives the network broadcast request forwards the request to every network node with which it has CP-CP sessions. (It does not forward the request to the node from which it received the broadcast request.) Those nodes then begin searching their respective domains for the target resource. (For broadcast-specific flows for those nodes, see Figure 198 on page 555.)
4. NN2 searches its domain for the resource. (Flows are not shown; see Figure 198 on page 555.) NN2 does not locate the resource. However, NN2 does not reply to NN1 until it has received a reply from all of the nodes to which it forwarded the request. NN3 does not locate the resource in its own domain and replies to NN2.
5. NN2 now returns a negative reply to NN1 because NN2 has exhausted its search logic.
6. NN4 owns the resource; therefore, it returns a positive reply to NN1.

7. Because the resource discovery search located the resource, NN1 sends a search to the target containing the original session-specific information.
8. The target is found.

Locate resource: EN to NN to subarea network

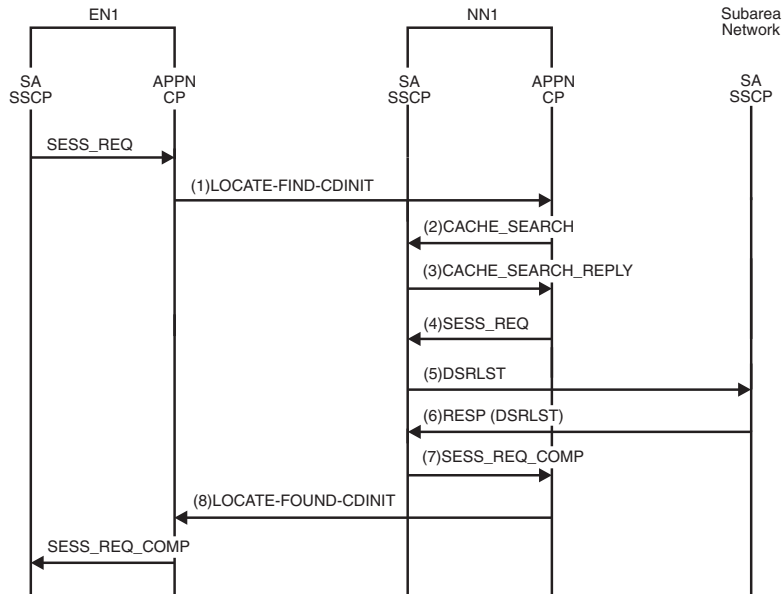


Figure 186. Locate resource: EN to NN to subarea network

1. The end node sends a search request for a target resource to the network node server. As network node server of the originating LU, NN1 looks for the target resource in the directory database. NN1 has no knowledge of the target resource in the directory database.
2. The APPN control point (CP) requests that the subarea SSCP check for information about the location of the resource.
3. The subarea SSCP replies that an entry for the target resource is found.
4. The APPN CP requests that the subarea SSCP send the search request to the target resource.
5. The request is sent to the owning SSCP.
6. The owning SSCP indicates that it owns the target resource.
7. The subarea SSCP replies to the APPN CP that the target resource is found.
8. NN1 returns a positive reply to EN1.

Locate resource: Complex APPN network

3. The subarea SSCP replies that the target resource is not known.
4. Because NN1 does not have the target resource in either its APPN or subarea directories, it initiates a resource discovery search for the resource. Because the resource discovery search starts at the beginning of the search logic, another CACHE_SEARCH is performed. NN1 has no APPN-domain end nodes; therefore, no domain broadcast occurs.
5. The subarea SSCP replies that the target resource is not known.
6. NN1 does not initiate a network broadcast because there is a central directory server (CDS) in the network; therefore, NN1 sends a request to this CDS.
7. CDS1 receives the request and performs origin CDS logic. CDS1 looks in its directory database for the target resource and has an entry that indicates that the target resource resides on NN4. CDS1 sends the request to NN4.
8. NN4 owns the resource; therefore, NN4 returns a positive reply.
9. CDS1 replies to NN1.
10. Because the resource discovery search found the resource, NN1 sends a search to the target, containing the original session-specific information.
11. The target is found.
12. NN1 replies to the end node.

Locate resource: Complex APPN network using more than one CDS

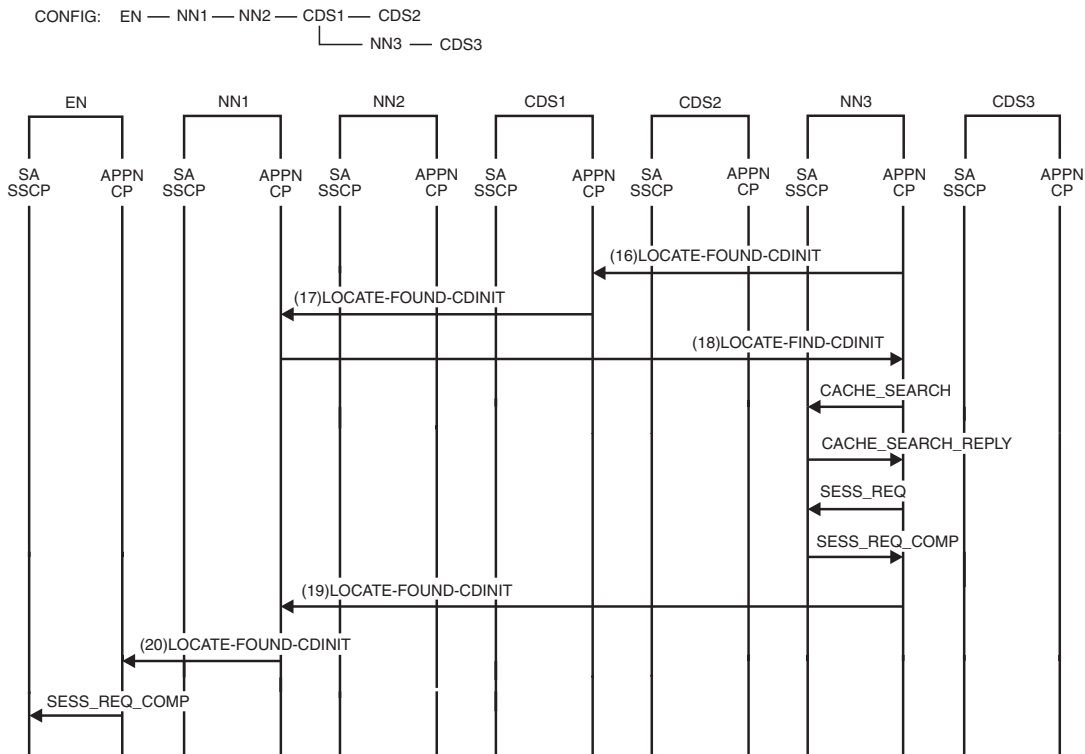


Figure 189. Locate resource: Complex APPN network using more than one CDS (part 2 of 2)

1. The end node sends a search request for a target resource to the network node server. As network node server of the originating LU, NN1 looks for the target resource in the directory database. NN1 has no knowledge of the target resource in the directory database.
2. The APPN control point (CP) requests that the subarea SSCP check its resource information for information about the location of the resource.
3. The subarea SSCP replies that the target resource is not known.
4. Because NN1 does not have the target resource in either its APPN or subarea directories, it initiates a resource discovery search for the resource. Because the resource discovery search starts at the beginning of the search logic, another CACHE_SEARCH is performed.
5. The subarea SSCP replies that the target resource is not known.
6. NN1 has no APPN-domain end nodes; therefore, no domain broadcast occurs. NN1 does not initiate a network broadcast because there is a central directory server (CDS) in the network. NN1 sends a request to CDS1.
7. CDS1 receives the request and performs origin CDS logic. CDS1 looks in its directory database for the target resource and does not have an entry. Therefore, the APPN control point (CP) requests the location of the resource from the subarea SSCP.
8. The subarea SSCP replies that the target resource is not known.
9. CDS1 has no domain end nodes; therefore, no domain broadcast occurs. CDS1 begins an alternate CDS search by sending a request to CDS2.
10. CDS2 looks in its directory database for the target resource and does not have an entry. The APPN CP requests that the subarea SSCP check for information about the location of the resource.

11. The subarea SSCP replies that the target resource is not known.
12. CDS2 has no domain end nodes to which to send a domain broadcast. Therefore, CDS2 returns a negative reply to CDS1.
13. CDS1 continues the alternate CDS search by sending a request to CDS3.
14. CDS3 looks in its directory database and finds an entry that indicates that the target resource resides on NN3. CDS3 replies to CDS1 with this information.
15. CDS1 sends a request to NN3 to verify that the target actually resides there.
16. NN3 replies to CDS1 that it does, indeed, own the target resource.
17. CDS1 returns a reply to NN1.
18. Because the resource discovery search found the resource, NN1 sends a search to the target, containing the original session-specific information.
19. The target is found.
20. NN1 returns a reply to the end node.

Locate resource: Network node server, NN1, of the originating logical unit (OLU) is at pre-V4R2 level

CONFIG: EN — NN1 — NN2 — CDS1

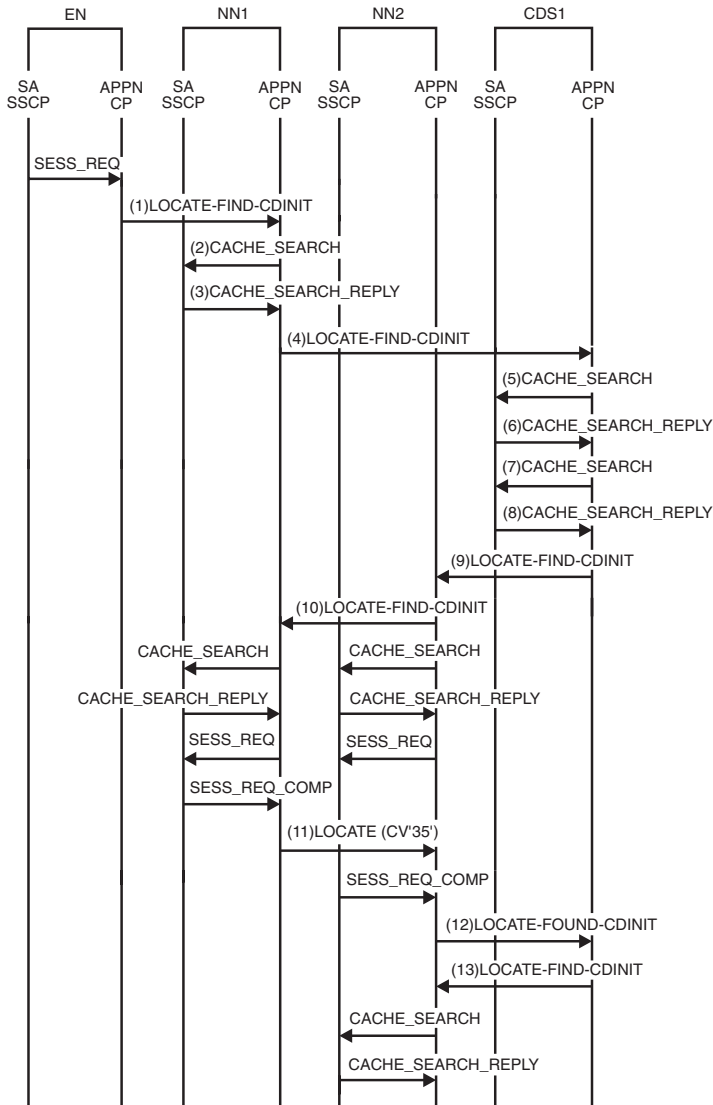


Figure 190. Locate resource: Network node server, NN1, of the originating logical unit (OLU) is at pre-V4R2 level (part 1 of 2)

CONFIG: EN — NN1 — NN2 — CDS1

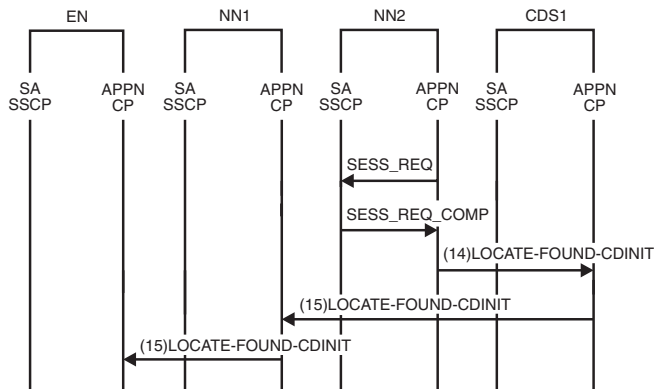


Figure 191. Locate resource: Network node server, NN1, of the originating logical unit (OLU) is at pre-V4R2 level (part 2 of 2)

1. The end node sends a search request for a target resource to the network node server. As network node server of the originating LU, NN1 looks for the target resource in the directory database. NN1 has no knowledge of the target resource in the directory database.
2. The APPN control point (CP) requests that the subarea SSCP check its resource information for the location of the resource.
3. The subarea SSCP replies that the target resource is not known. NN1 has no APPN end nodes; therefore, no domain broadcast occurs.
4. Because there is a central directory server (CDS) in the network, NN1 does not initiate a network broadcast. Instead, NN1 sends a request to CDS1.
5. The CDS does not know the target resource in its directory. The APPN control point (CP) requests that the subarea SSCP check its resource information for the location of the resource.
6. The subarea SSCP replies that the target resource is not known.
7. Because CDS1 does not have the target resource in either its APPN or subarea directories, it initiates a resource discovery search. Because the resource discovery search starts at the beginning of the search logic, another CACHE_SEARCH is performed.
8. The subarea SSCP replies that the target resource is not known.
9. CDS1 has no APPN end nodes; therefore, no domain broadcast occurs. CDS1 initiates a network broadcast for the target resource.
10. NN2 forwards the network broadcast request and then begins to search its domain.
11. After completing its search logic, NN1 returns a negative reply to the network broadcast request.
12. The subarea SSCP on NN2 indicates that the resource is found; therefore, a positive reply is returned.
13. Because the broadcast that was initiated by the resource discovery search found the resource, the original search request containing session-specific information is sent to the target location, NN2.
14. NN2, the owner of the resource, returns a positive reply.
15. CDS1 replies to NN1, and NN1 replies to the end node.

Locate resource: APPN and subarea network

CONFIG: EN—NN1—CDS1—ICN2==Subarea Network==ICN3—NN4

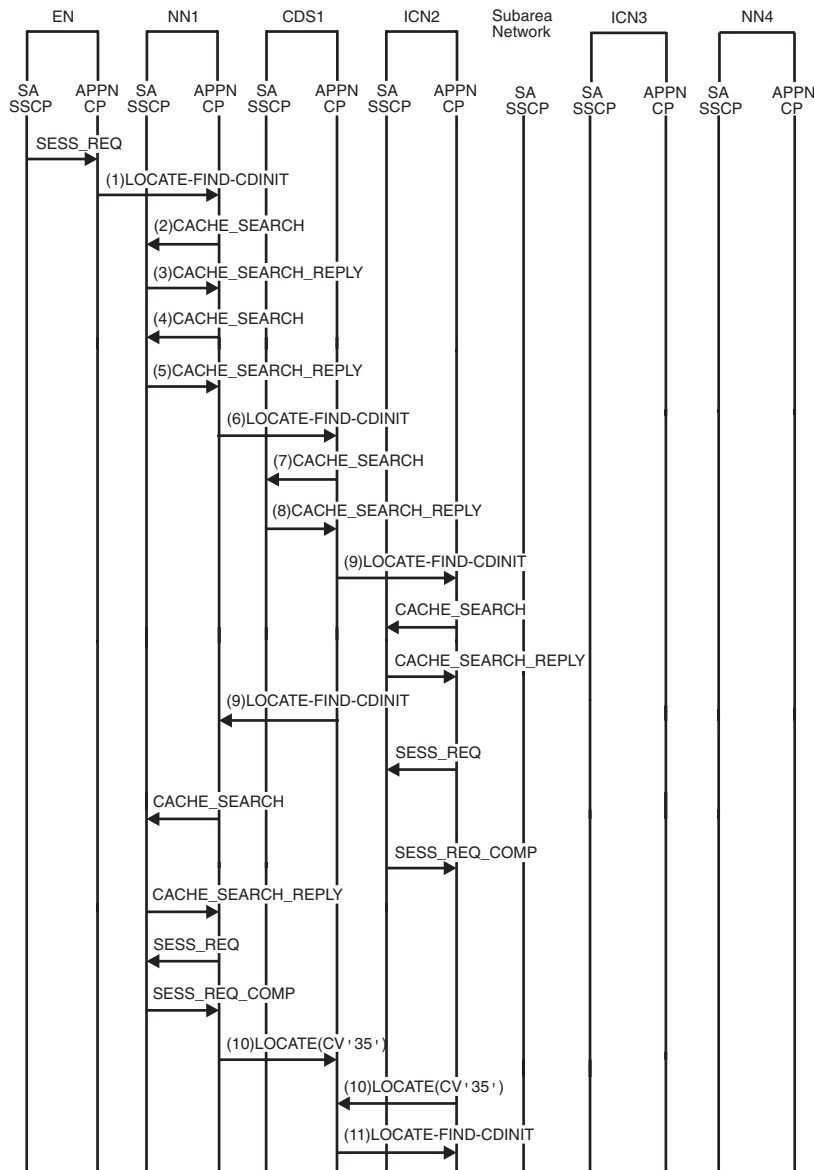


Figure 192. Locate resource: APPN and subarea network (part 1 of 3)

CONFIG: EN — NN1 — CDS1 — ICN2==Subarea Network==ICN3 — NN4

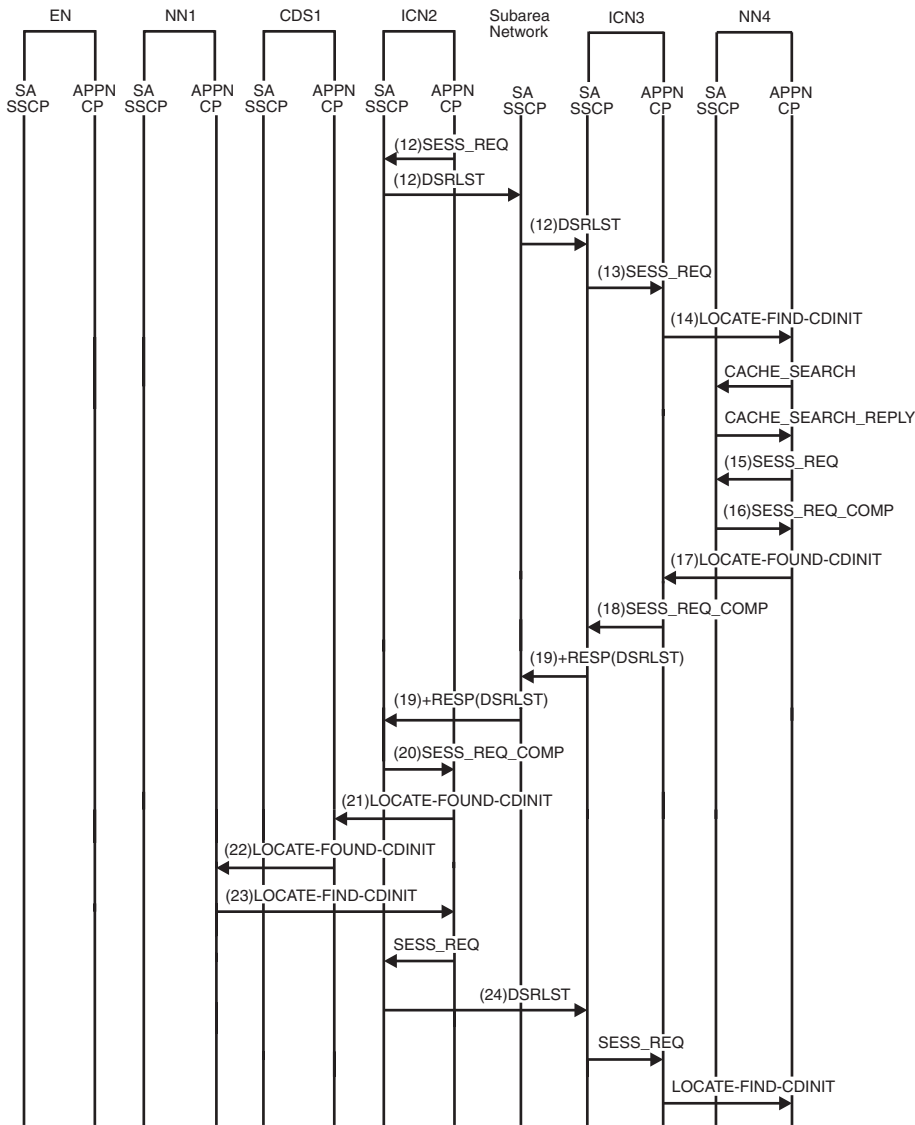


Figure 193. Locate resource: APPN and subarea network (part 2 of 3)

CONFIG: EN—NN1—CDS1—ICN2==Subarea Network==ICN3—NN4

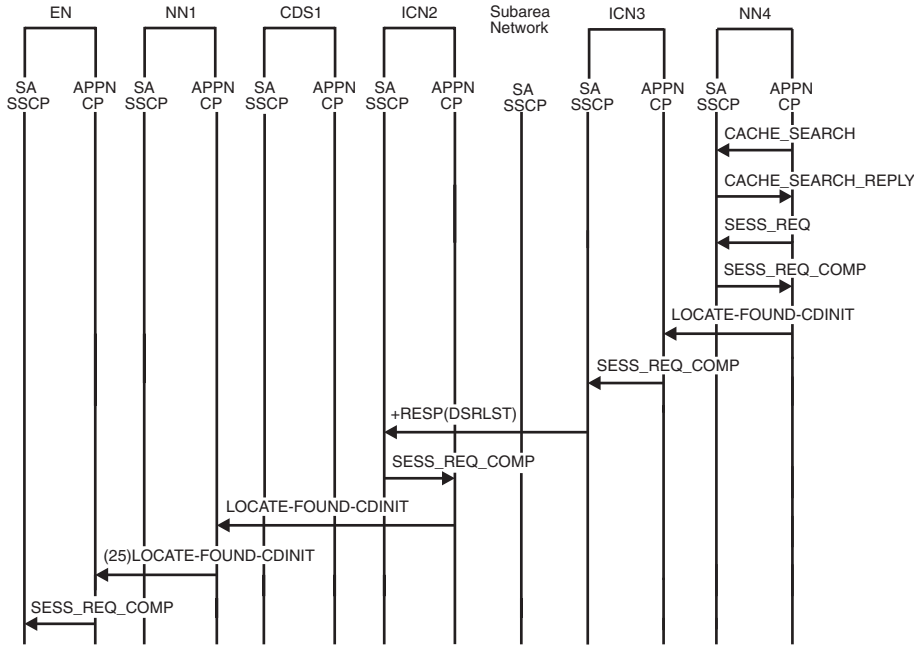


Figure 194. Locate resource: APPN and subarea network (part 3 of 3)

1. The end node sends a search request for a target resource to the network node server. As network node server of the originating LU, NN1 looks for the target resource in the directory database. NN1 has no knowledge of the target resource in the directory database.
2. The APPN control point (CP) requests that the subarea SSCP check for information about the location of the resource.
3. The subarea SSCP replies that the target resource is not known.
4. Because NN1 does not have the target resource in either its APPN or subarea directories, it initiates a resource discovery search for the resource. Because the resource discovery search starts at the beginning of the search logic, another CACHE_SEARCH is performed.
5. The subarea SSCP replies that the target resource is not known.
6. NN1 has no APPN-domain end nodes; therefore, no domain broadcast occurs. NN1 does not initiate a network broadcast because there is a central directory server (CDS) in the network; therefore, NN1 sends a request to CDS1.
7. CDS1 receives the request and performs origin CDS logic. CDS1 looks in its directory database for the target resource and does not have an entry. Therefore, the APPN control point (CP) requests the location of the resource from the subarea SSCP.
8. The subarea SSCP replies that the target resource is not known.
9. CDS1 has no domain end nodes; therefore, no domain broadcast occurs. There are no other CDSs in the network; therefore, no alternate CDS search occurs. CDS1 initiates a network broadcast. CDS1 sends the broadcast request to all nodes with which it has CP-CP sessions.

Note:

- a. CDS1 must send the broadcast request to NN1, even though NN1 originated the request to CDS1, because there can be parts of the APPN network that are reachable only through NN1 (these parts are not shown here).
 - b. The network broadcast sent by a CDS indicates that attached subarea networks should not be searched at this time.
10. Both NN1 and ICN2 respond that the target is not found.
 11. After CDS1 has collected all the replies from the network broadcast, CDS1 continues the search with an interchange node search. CDS1 sends the interchange node search request to ICN2. This request indicates that the interchange node is to search its attached subarea network.
 12. The APPN CP requests that the subarea SSCP initiate subarea routing. ICN2 sends DSRLST to adjacent SSCP.
 13. The subarea SSCP requests that the APPN CP initiate APPN searching.
 14. ICN3 looks for the target resource in the directory database. ICN3 has knowledge in its directory database that the target resource resides in NN4. ICN3 sends a directed search request to NN4.
 15. The APPN CP sends a SESS_REQ signal to the subarea SSCP.
 16. The subarea SSCP sends a SESS_REQ_COMP to the APPN CP indicating that the target resource is located.
 17. NN4 returns a positive reply to NN3.
 18. The APPN CP replies to the subarea SSCP.
 19. Positive responses to DSRLSTs are returned.
 20. The subarea SSCP replies to the APPN CP.
 21. ICN2 replies to CDS1.
 22. CDS1 replies to NN1.
 23. Because the resource discovery search located the resource, NN1 sends a search to the target, containing the original session-specific information.
 24. A DSRLST, which contains session-specific information, is sent.
 25. NN1 replies to the end node.

Locate resource: Directory search verification reduction

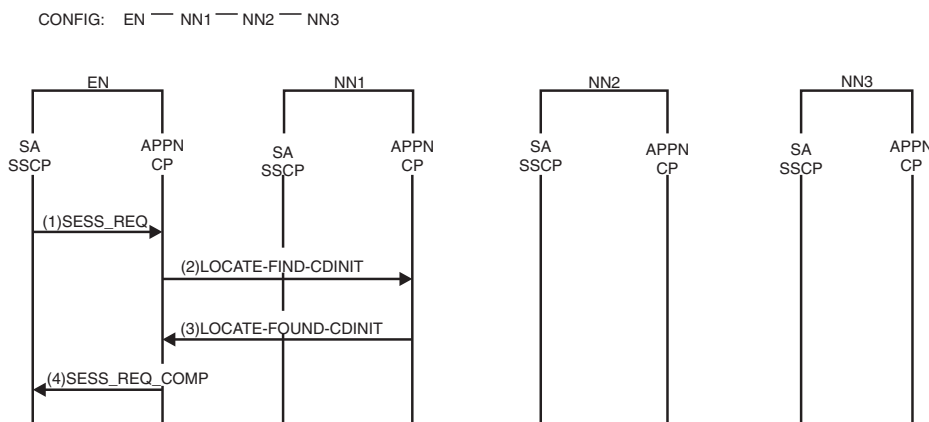


Figure 195. Locate resource: Directory search verification reduction

1. The subarea SSCP requests information about a resource and does not require the APPN CP to verify the location of the resource.
2. The end node sends a search request for the target resource to its network node server. Only network node servers, when requested by the origin node, can respond to search requests without first verifying the location of the resource.
3. NN1 looks for the target resource in its directory database. NN1 has knowledge in its directory that the target resource resides on NN3. Further, NN1 has information that allows the search to succeed, without verifying the location of the target resource. Therefore, NN1 returns a positive reply on behalf of NN3.
4. The APPN CP replies to the subarea SSCP.

Locate resource: SLU-initiated session

CONFIG: SLU EN1—NN1—NN2—NN3—EN2

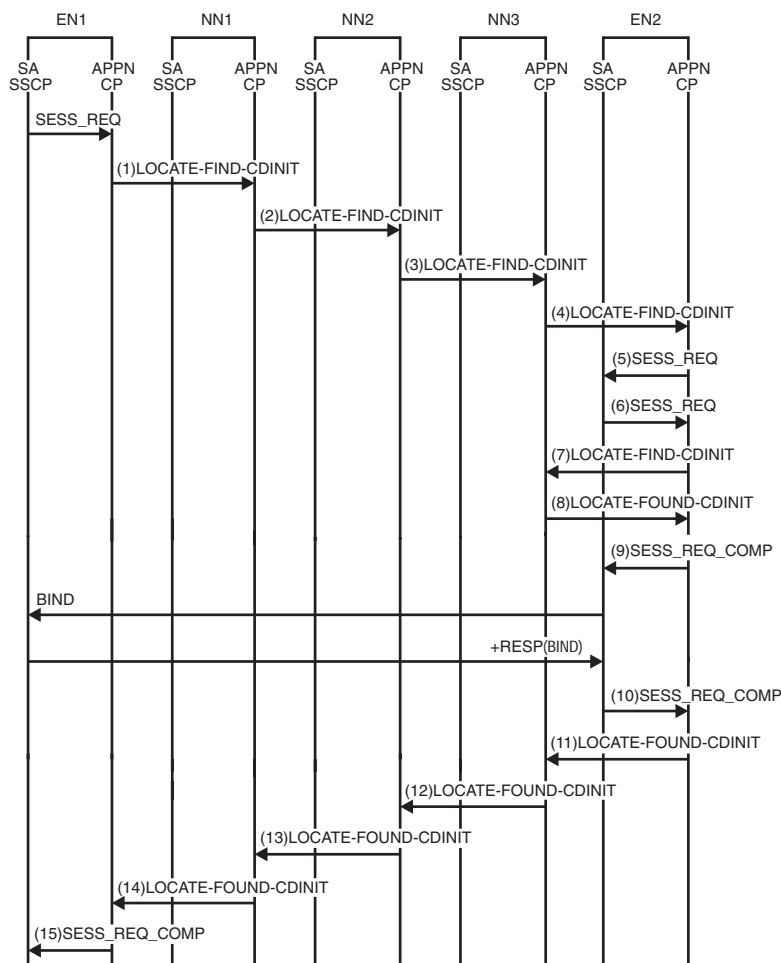


Figure 196. Locate resource: SLU-initiated session

1. The end node sends a session request for a primary LU (PLU) to its network node server. NN1 looks for the PLU in its directory database. NN1 has knowledge in the directory database that the PLU resides on EN2, which is served by NN3.

2. The network node sends a directed search request to NN3. Because NN1 does not have direct CP-CP sessions with NN3, NN1 sends the directed search request to NN3 through NN2.
3. NN2 is not the destination of the directed search; therefore, NN2 forwards the request to NN3.
4. NN3 receives the request and forwards it to EN2.
5. The APPN CP sends a SESS_REQ signal to the subarea SSCP.
6. The PLU initiates a search for the secondary LU (SLU), indicating that the location of the target does not have to be verified but that an RSCV must be calculated.
7. The end node sends a search request for the SLU to its network node server.
8. NN3 looks for the SLU in its directory database. NN3 has knowledge in its directory database that the SLU resides on EN1. Further, NN3 has information that allows the search to succeed without verifying the location of the SLU. On behalf of EN1, NN3 returns a positive reply, which includes the RSCV for the session.
9. The APPN CP replies to the subarea SSCP for the PLU-initiated request.
10. The subarea SSCP replies to the APPN CP for the SLU-initiated request, indicating that the session is already active.
11. The end node sends a reply for the SLU-initiated request to NN3.
12. NN3 replies to NN2.
13. NN2 replies to NN1.
14. NN1 replies to EN1.
15. The APPN CP replies to the subarea SSCP.

Locate resource: CP-CP session terminates

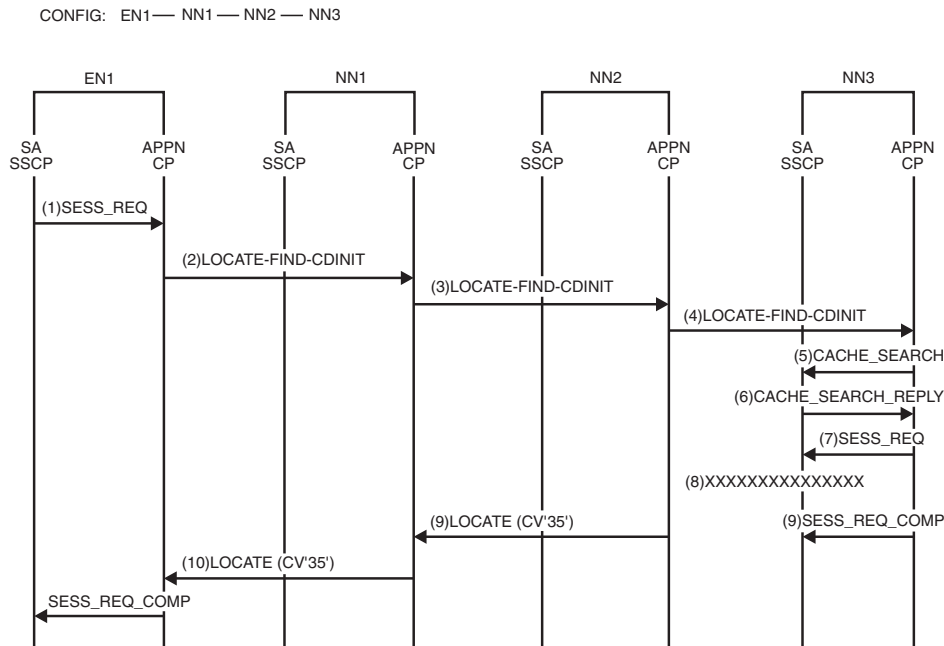


Figure 197. Locate resource: CP-CP session terminates

1. The subarea SSCP requests a search for a target resource.
2. The end node sends a search request for a target resource to the network node server. The network node looks for the target resource in the directory database. The network node has knowledge in the directory that the target resource resides on NN3.
3. The network node sends a directed search request to NN3. Because NN1 does not have direct CP-CP sessions with NN3, NN1 sends the directed search request to NN3 through NN2.
4. NN2 is not the destination of the directed search; therefore, NN2 forwards the request to NN3.
5. The APPN control point (CP) requests that the subarea SSCP check its resource information for information about the location of the resource.
6. The subarea SSCP replies that the target resource is found.
7. The APPN CP requests that the subarea SSCP initiate a search for the target resource.
8. The CP-CP session goes down between NN2 and NN3.
9. NN2 sends a negative reply to NN1. NN3 cleans up its control blocks.
10. NN1 continues its search logic. If another path to the target exists, the target can be found (for example, through a network broadcast search). Otherwise, the search does not find the target resource.

Locate resource: Network node receives network broadcast request

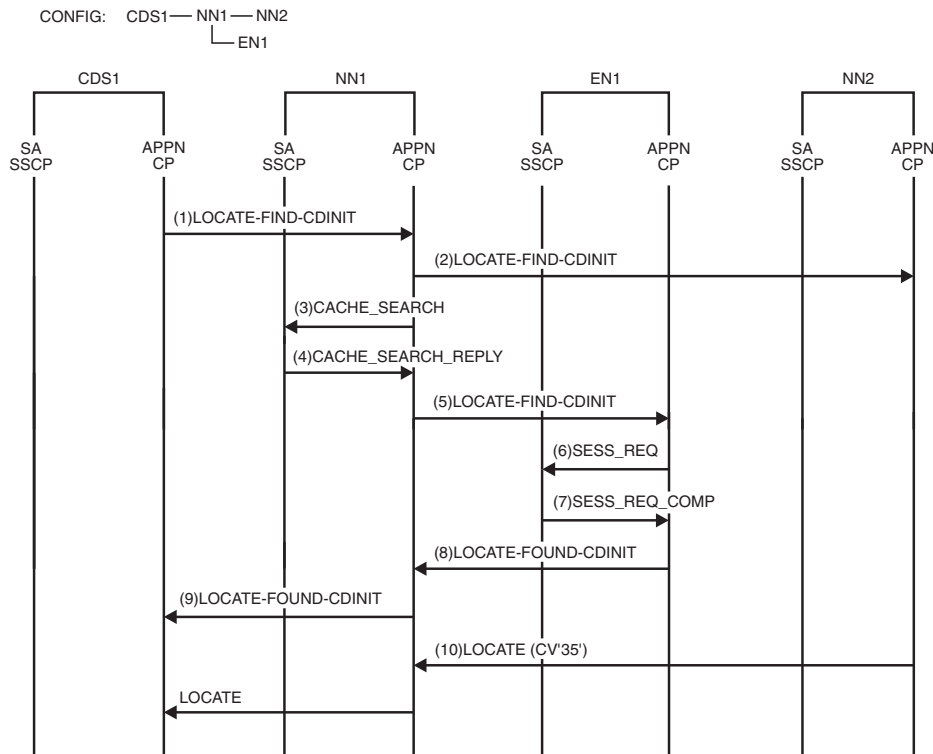


Figure 198. Locate resource: Network node receives network broadcast request

Note: This figure illustrates only the flows at NN1.

1. CDS1 has reached the point in its search logic where a network broadcast is performed. CDS1 sends the network broadcast to all network nodes with which it has CP-CP sessions.
2. NN1 recognizes that the request is a network broadcast. Therefore, it forwards the request immediately to all network nodes with which it has CP-CP sessions.
3. NN1 then begins to search itself and its domain. NN1 looks for the target resource in its directory database. NN1 does not have an entry for the target resource. The search continues with the APPN CP requesting that the subarea SSCP check its resource information for information about the location of the resource.
4. The subarea SSCP replies that the target is not known.
5. NN1 sends the domain broadcast request to all served ENs that indicated in the CP_CAPABILITIES exchange that they are to be searched on domain broadcasts.
6. The APPN CP forwards the search to the subarea SSCP.
7. The subarea SSCP indicates that the target resource was found.
8. EN1 returns a positive reply.
9. NN1 forwards the positive reply to CDS1 and continues to wait for a reply from NN2. NN1 indicates to CDS1 in its reply that NN1 has not yet received replies to all of its searches. CDS1 forwards the positive reply to the originator of the search and continues to wait for a final reply from NN1.
10. NN2 returns a negative reply to NN1. Because NN1 has now completed all of its searches, NN1 returns a Neutral reply to CDS1 to indicate that NN1 has

completed its search. A Neutral reply is one that contains neither a CV'35' (to indicate failure) or a FOUND GDS variable (to indicate success). A CDS receiving a neutral reply does not forward the neutral reply to the originator of the search. A CDS returns only one reply to the originator to indicate success or failure.

LU-LU session flows

Figure 199 on page 558 through Figure 237 on page 587 show the network flows to establish LU-LU sessions.

The figure captions for some of the figures indicate the configuration that the flow applies to and the type of session shown. For example, the caption EN (PLU)–NNS...APPN Network; PLU-Initiated, with No Queueing means *A primary logical unit (PLU) on an end node (EN) that is attached to an APPN network through a network node server (NNS)*. The symbol — indicates a CP-CP session. The symbol == indicates a CDRM-CDRM session. The symbol ... means that part of the network is not shown.

All of these flows assume that the directory services database has accurate information about the location of the destination LU.

The following terms are used in these figures:

Term Meaning

Endpoint TGVs

A list of control vector pairs: CV X'46' and X'47'

RSCV Route selection control vector, CV X'2B'

Scout search

Sent to find out the location of the destination LU (DLU) and to precompute the session RSCV, if either the origin LU (OLU) or the DLU is in a subarea network. Because it is necessary only to find the DLU and not to set up the session, it is not necessary to actually verify the location of the DLU or to reserve resources for the session.

Index of LU-LU session flows

Table 42 lists the LU-LU session flows illustrated here.

Table 42. Index of LU-LU session flows

Flow	Page
APPN network...NNS--EN (PLU)	
SLU-initiated, no queueing	Figure 206 on page 562
SLU-initiated, queued by the PLU	Figure 207 on page 562
APPN network...NNS--EN (SLU)	
PLU-initiated, no queueing	Figure 204 on page 561
PLU-initiated, queued by the SLU	Figure 205 on page 561
APPN network (PLU)...ICN==SA(SLU), PLU-Initiated	
Directed search without required precomputed RSCV	Figure 230 on page 580
No queueing	Figure 229 on page 579

Table 42. Index of LU-LU session flows (continued)

Flow	Page
Queued by SLU	Figure 232 on page 582
Search-only flow transformed into a DSRLST	Figure 228 on page 578
USERVAR resolution required	Figure 231 on page 581
APPN network (PLU)...ICN==VR-based TG==ICN...APPN network (SLU)	
PLU-initiated	Figure 237 on page 587
APPN network (SLU)...ICN==PLU	
Autologon, PLU not available initially	Figure 236 on page 586
SLU-initiated, no queueing	Figure 233 on page 583
SLU-initiated, queued by the PLU	Figure 234 on page 584
CLSDST PASS; SLU is single-session capable	
From APPN to subarea	Figure 225 on page 576
Through APPN	Figure 224 on page 575
EN-NN-EN, PLU-initiated, no queueing (including BIND flows for intermediate network node)	Figure 226 on page 577
EN (PLU)--NNS...APPN network	
PLU-initiated, no queueing	Figure 199 on page 558
PLU-initiated, queued by the PLU	Figure 200 on page 559
PLU-initiated, queued by the SLU	Figure 201 on page 559
EN (SLU)--NNS...APPN network	
SLU-initiated, no queueing	Figure 202 on page 560
SLU-initiated, queued by the PLU	Figure 203 on page 560
Intermediate Network Node (INN) BIND. The LOCATE did not go through this node.	Figure 227 on page 578
SA (PLU)==ICN...APPN network (SLU)	
DSRLIST transforming into PLU-initiated, search-only	Figure 208 on page 563
PLU-initiated, no queueing	Figure 209 on page 564
PLU-initiated, queued by the SLU	Figure 211 on page 565
PLU-initiated, USERVAR resolution required	Figure 210 on page 565
SA (SLU)==ICN...APPN network (PLU)	
Autologon, PLU not available initially	Figure 215 on page 569
SLU-initiated, no queueing	Figure 213 on page 567
SLU-initiated, queued by the PLU	Figure 214 on page 568
Session release request	
SA(PLU)==ICN...APPN network(SLU)	Figure 222 on page 574
SA(SLU)==ICN...APPN network(PLU)	Figure 223 on page 574
Session termination, forced	
SA(PLU)==ICN...APPN Network(SLU), pending active session. PLU is accessible without going into APPN.	Figure 219 on page 572
SA(PLU)==ICN...APPN Network(SLU), queued session	Figure 220 on page 573

Table 42. Index of LU-LU session flows (continued)

Flow	Page
SA(PLU)==ICN...APPN Network(SLU), queued session. PLU is accessible without going into APPN.	Figure 221 on page 573
SA(SLU)==ICN...APPN Network(PLU), pending active session	Figure 218 on page 571
Session termination, orderly	
SA(PLU)==ICN...APPN Network(SLU), active session	Figure 216 on page 570
APPN Network(PLU)...ICN==SA(SLU), active session	Figure 217 on page 571

EN (PLU)--NNS...APPN network, PLU-initiated, with no queuing

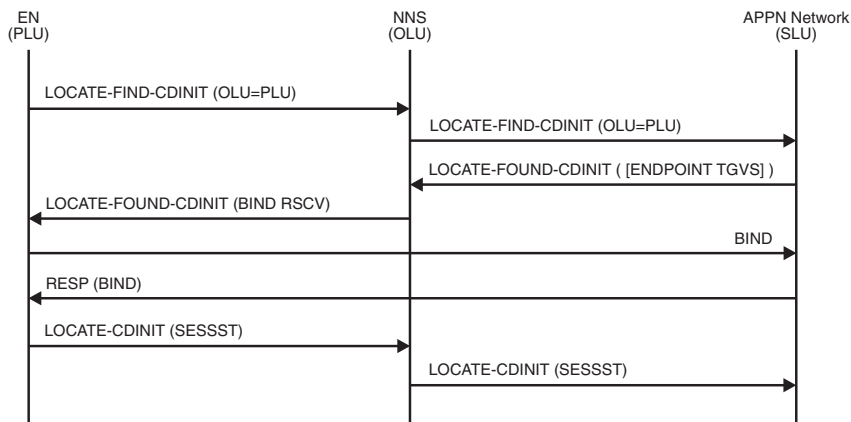


Figure 199. EN (PLU)—NNS...APPN network, PLU-initiated, with no queuing

Note: The BIND does not have to take the same path as the LOCATE flow.

EN (PLU)--NNS...APPN network, PLU-initiated, queued by the PLU

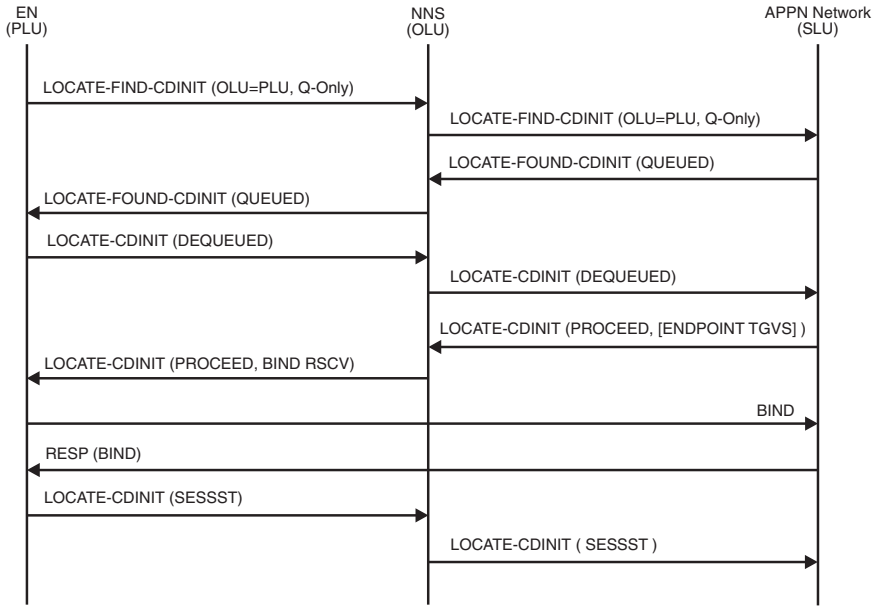


Figure 200. EN (PLU)—NNS...APPN network, PLU-initiated, queued by the PLU

Note: The BIND does not have to take the same path as the LOCATE flow.

EN (PLU)--NNS...APPN network, PLU-initiated, queued by the SLU

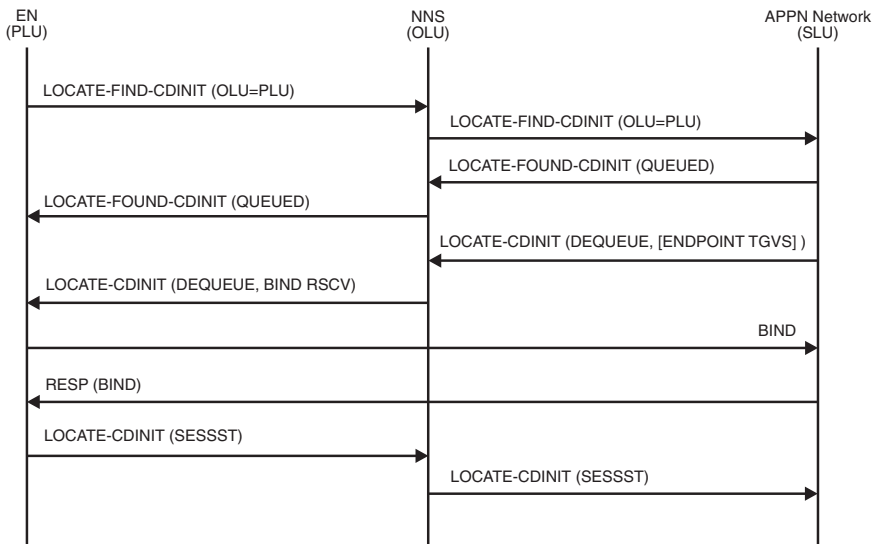


Figure 201. EN (PLU)—NNS...APPN network, PLU-initiated, queued by the SLU

Note: The BIND does not have to take the same path as the LOCATE flow.

EN (SLU)--NNS...APPN network, SLU-initiated, with no queuing

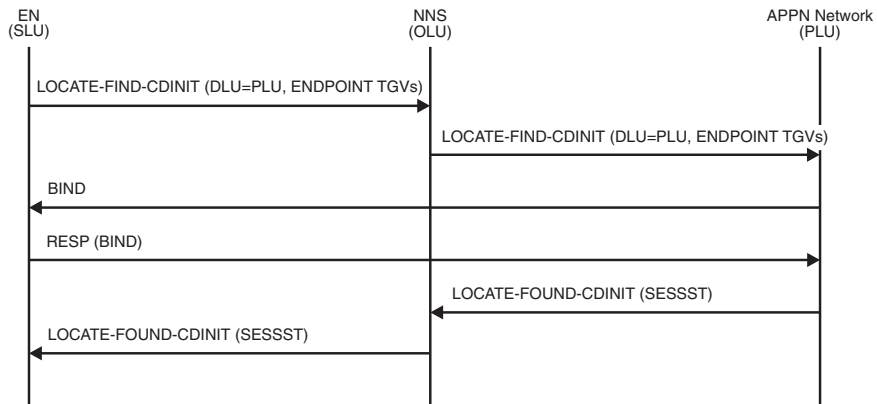


Figure 202. EN (SLU)—NNS...APPN network, SLU-initiated, with no queuing

Note: The BIND does not have to take the same path as the LOCATE flow.

EN (SLU)--NNS...APPN network, SLU-initiated, queued by the PLU

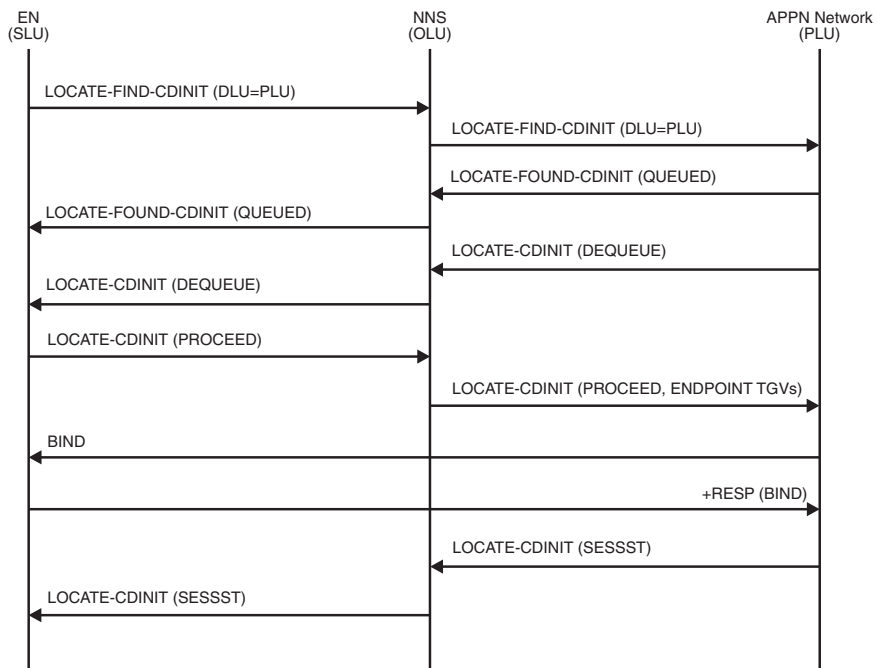


Figure 203. EN (SLU)—NNS...APPN network, SLU-initiated, queued by the PLU

Note: The BIND does not have to take the same path as the LOCATE flows.

APPN network...NNS--EN (SLU), PLU-initiated, no queuing

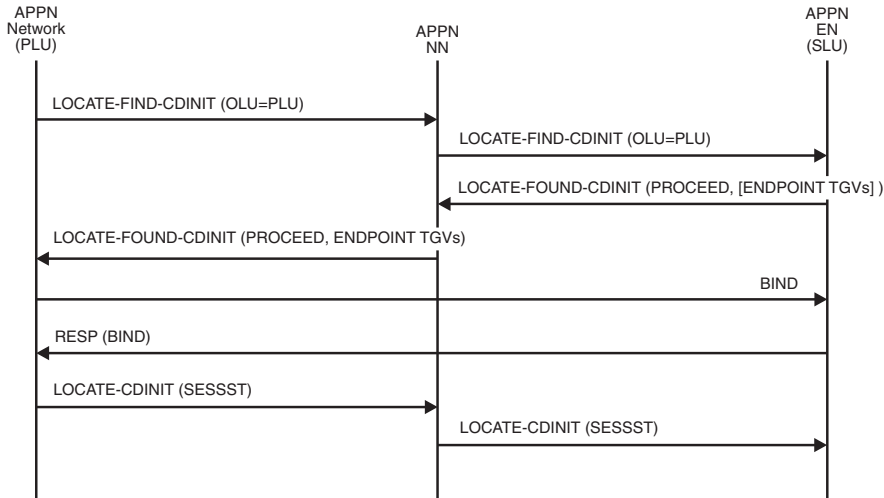


Figure 204. APPN network...NNS--EN (SLU), PLU-initiated, no queueing

Note: The BIND does not have to follow the same path as the LOCATE flow.

APPN network...NNS--EN (SLU), PLU-initiated, queued by the SLU

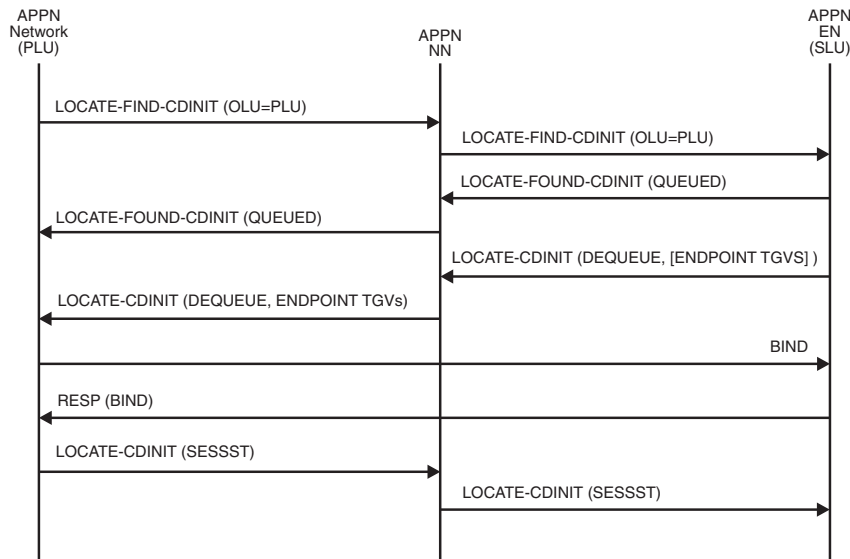


Figure 205. APPN network...NNS--EN (SLU), PLU-initiated, queued by the SLU

Note: The BIND does not have to follow the same path as the LOCATE flow.

APPN network...NNS--EN (PLU), SLU-initiated, no queueing

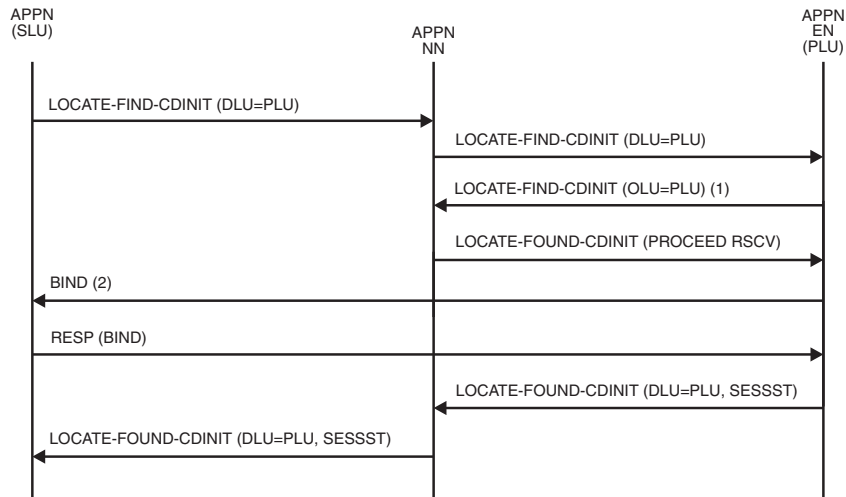


Figure 206. APPN network...NNS--EN (PLU), SLU-initiated, no queueing

1. The purpose of this search is to get the NN to compute the session RSCV for the EN.
2. The BIND does not have to follow the same path as LOCATE flows.

APPN network...NNS--EN (PLU), SLU-initiated, queued by the PLU

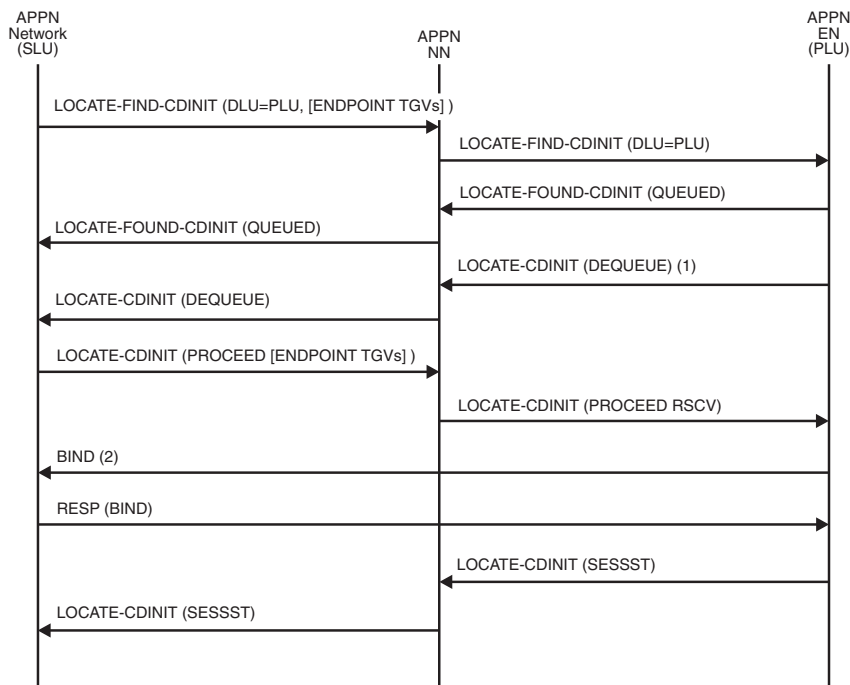


Figure 207. APPN network...NNS--EN (PLU), SLU-initiated, queued by the PLU

1. The PLU has become available.
2. The BIND does not have to follow the same path as the LOCATE flows.

SA (PLU)==ICN...APPN network (SLU), DSRLIST transforming into PLU-initiated, search-only

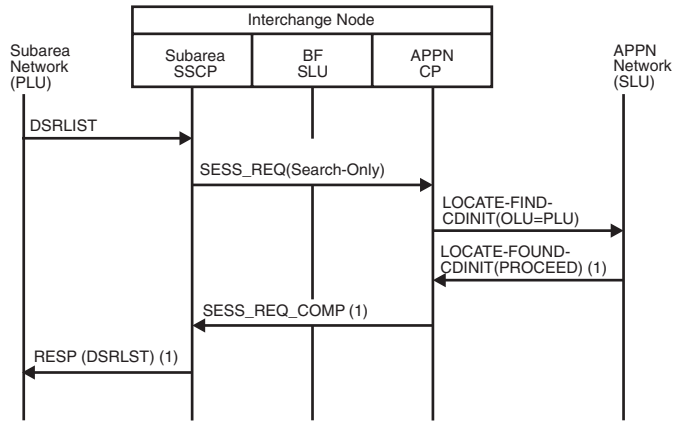


Figure 208. SA(PLU)==ICN...APPN network (SLU), DSRLIST transforming into PLU-initiated, search-only

1. Target LU location information.

SA (PLU)==ICN...APPN network (SLU), PLU-initiated, with no queuing

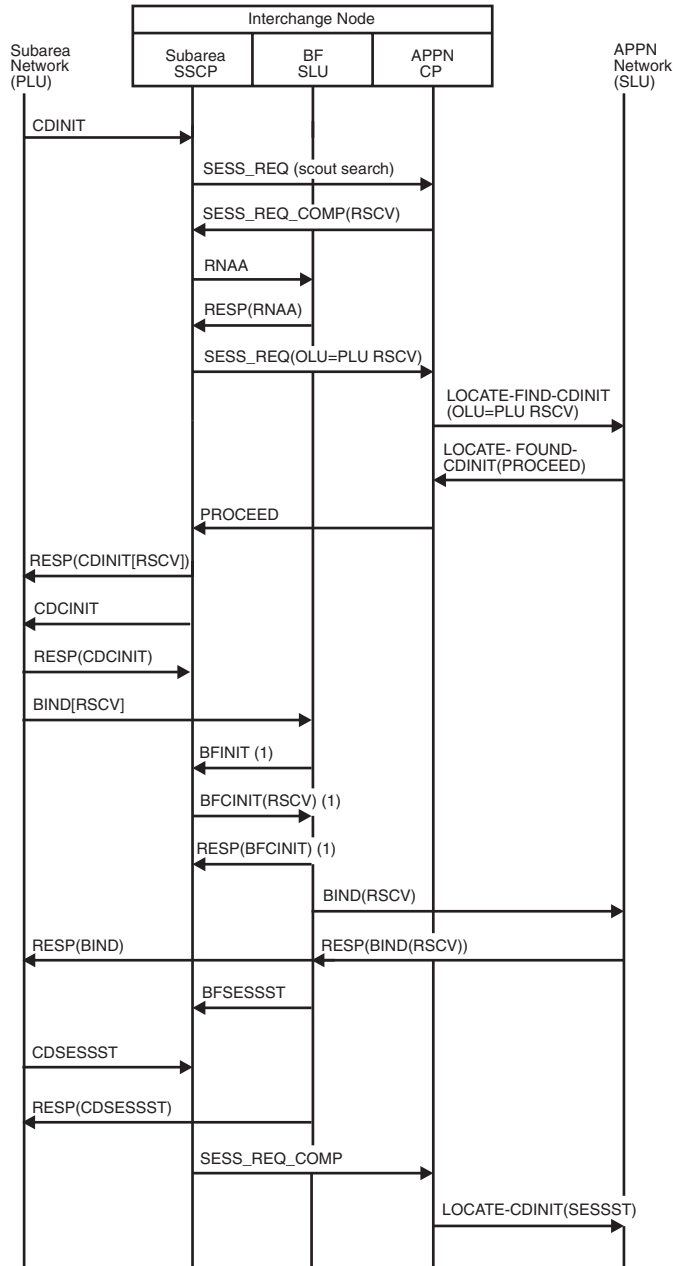


Figure 209. SA (PLU)==ICN...APPN network (SLU), PLU-initiated, with no queuing

1. These BFINIT/BFCINIT flows will not occur if the RSCV is passed on the BIND.

SA (PLU)==ICN...APPN network (SLU), PLU-initiated, USERVAR resolution required

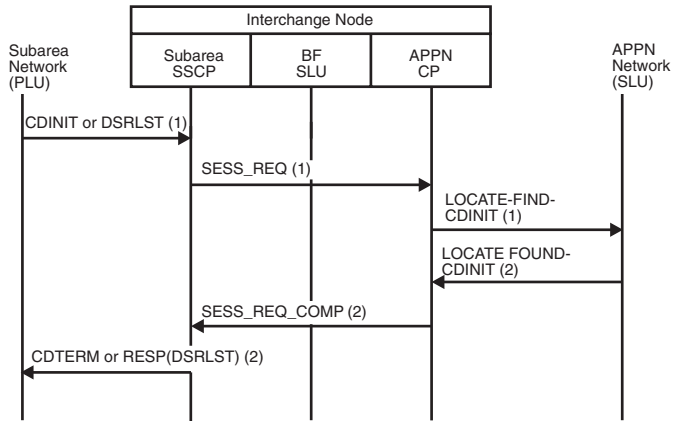


Figure 210. SA (PLU)==ICN...APPN network (SLU), PLU-initiated, USERVAR resolution required

1. Generic USERVAR name
2. Resolved USERVAR name

Note: For remaining session setup flows, see Figure 209 on page 564.

SA (PLU)==ICN...APPN network (SLU), PLU-initiated, queued by the SLU

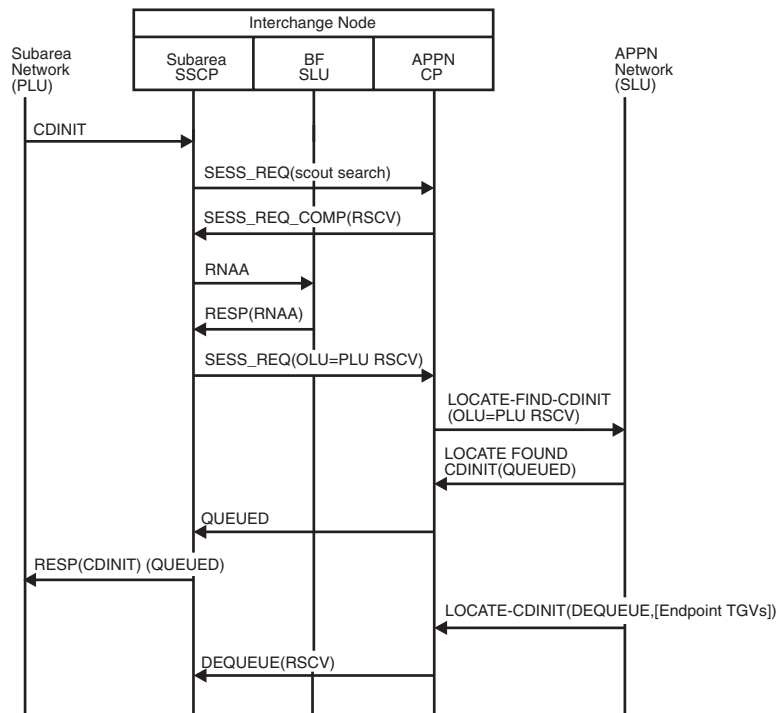


Figure 211. SA (PLU)==ICN...APPN network (SLU), PLU-initiated, queued by the SLU (part 1 of 2)

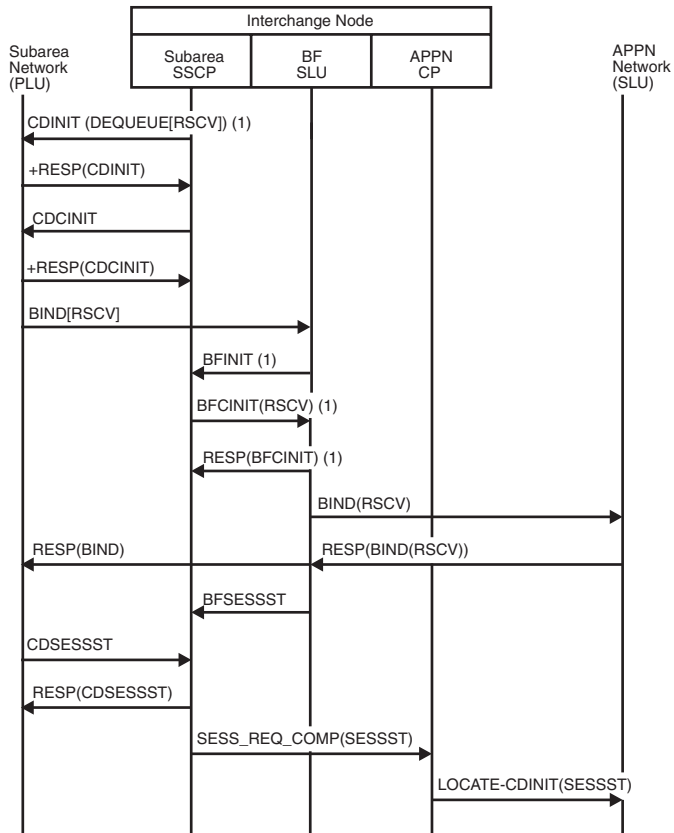


Figure 212. SA (PLU)==ICN...APPN network (SLU), PLU-initiated, queued by the SLU (part 2 of 2)

1. These BFINIT/BFCINIT flows will not occur if the RSCV is passed on the BIND.

SA (SLU)==ICN...APPN network (PLU), SLU-initiated, no queueing

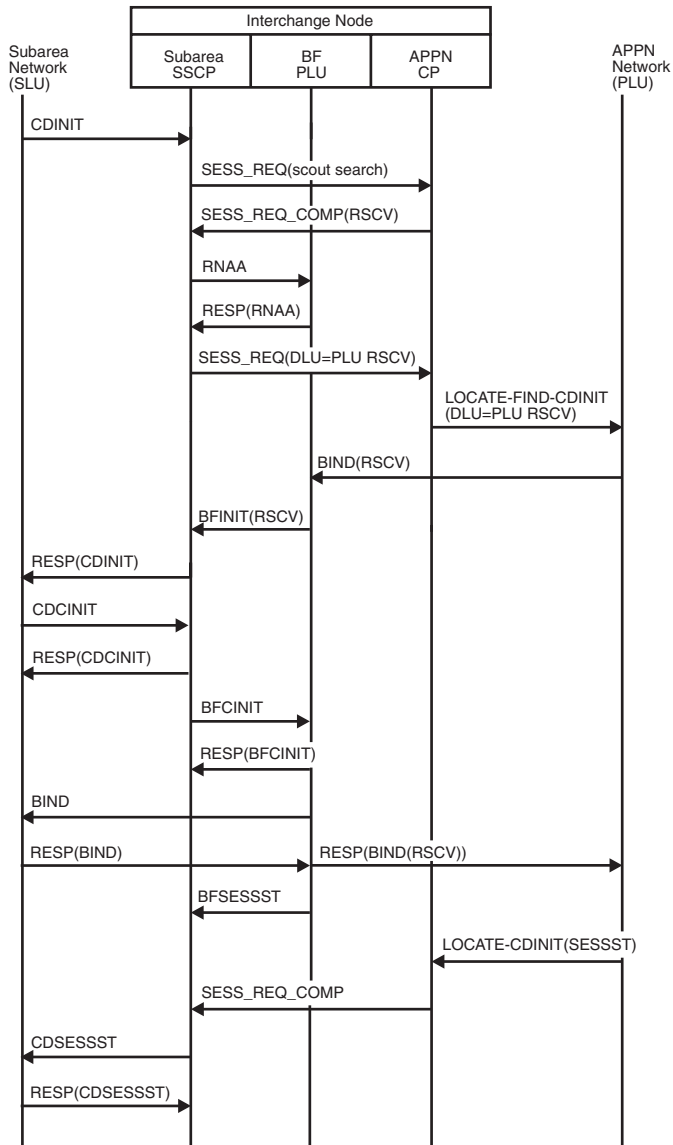


Figure 213. SA (SLU)==ICN...APPN network (PLU), SLU-initiated, no queueing

SA (SLU)==ICN...APPN network (PLU), SLU-initiated, queued by the PLU

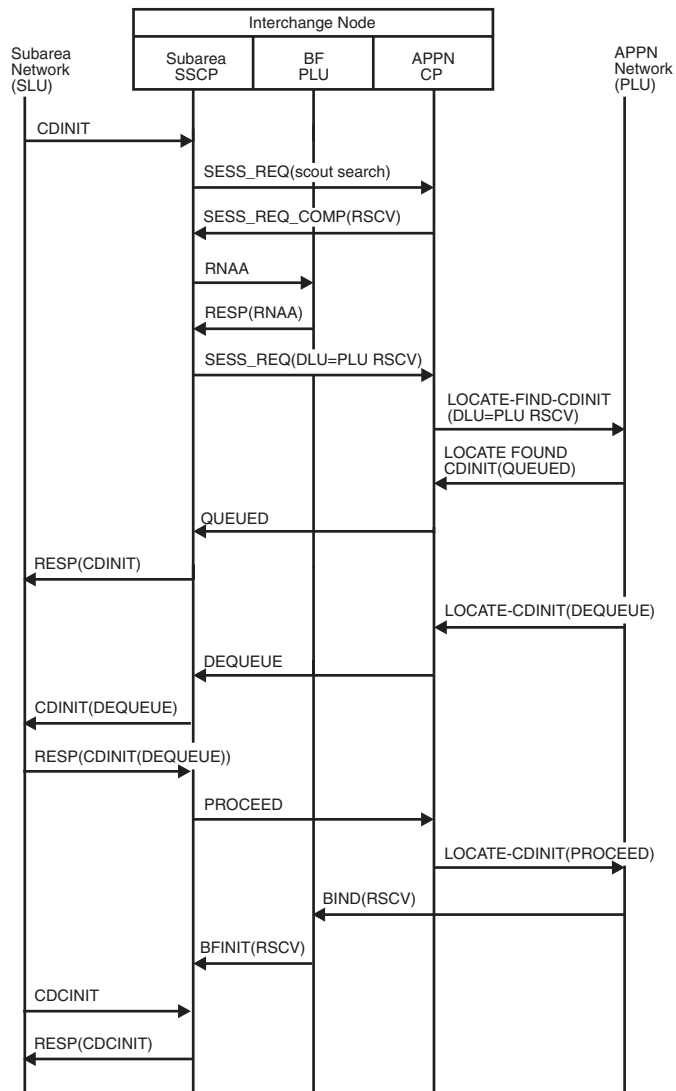


Figure 214. SA (SLU)==ICN...APPN network (PLU), SLU-initiated, queued by the PLU

SA (SLU)==ICN...APPN network (PLU), autologon, PLU not available initially

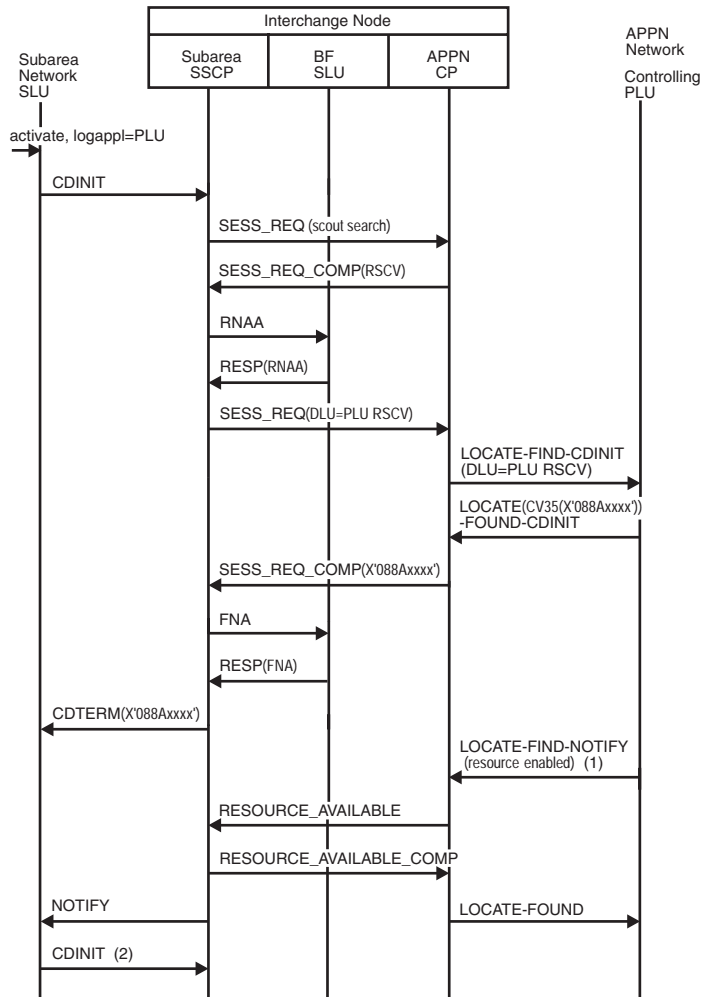


Figure 215. SA (SLU)==ICN...APPN network (PLU), autologon, PLU not available initially

1. The controlling PLU becomes available.
2. Normal SLU-initiated flows continue from here.

SA(PLU)==ICN...APPN network(SLU), orderly termination of active session

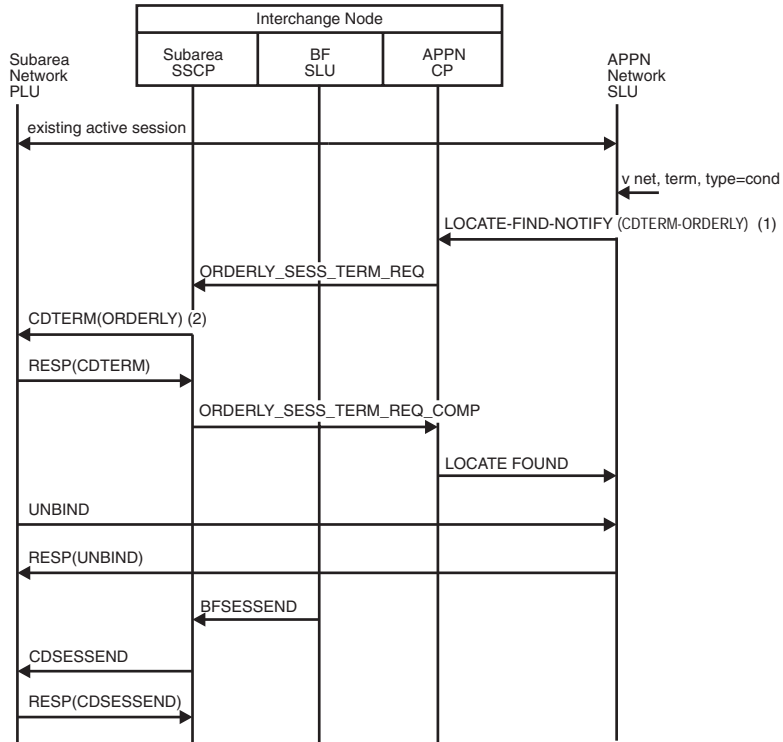


Figure 216. SA(PLU)==ICN...APPN network(SLU), orderly termination of active session

1. FQCPID of the session to be terminated.
2. The CDTERM type depends on the V NET,TERM type.

V NET,TERM type
CDTERM type

COND
 ORDERLY

UNCOND
 FORCED

FORCE
 CLEANUP

APPN network (PLU)...ICN==(SA)SLU, orderly termination of active session

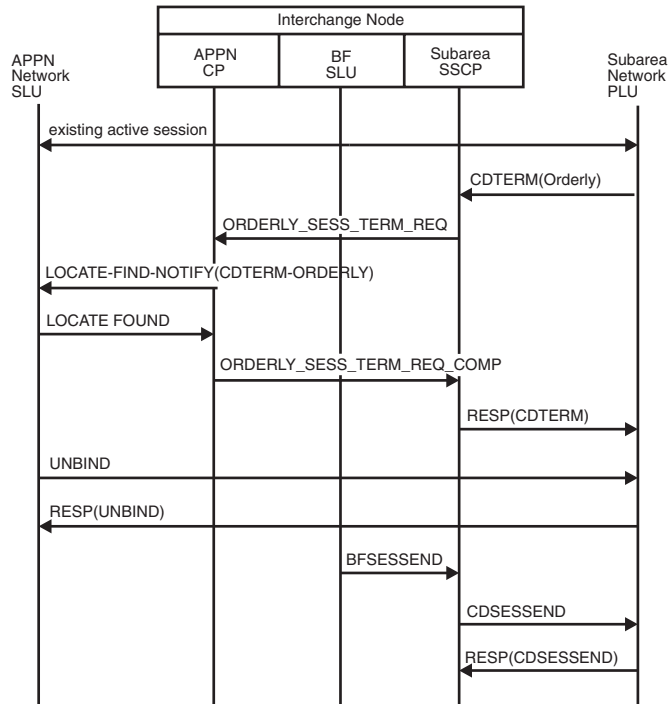


Figure 217. APPN network(PLU)...ICN==(SA)SLU, orderly termination of active session

SA(SLU)==ICN...APPN network(PLU), forced termination of pending active session

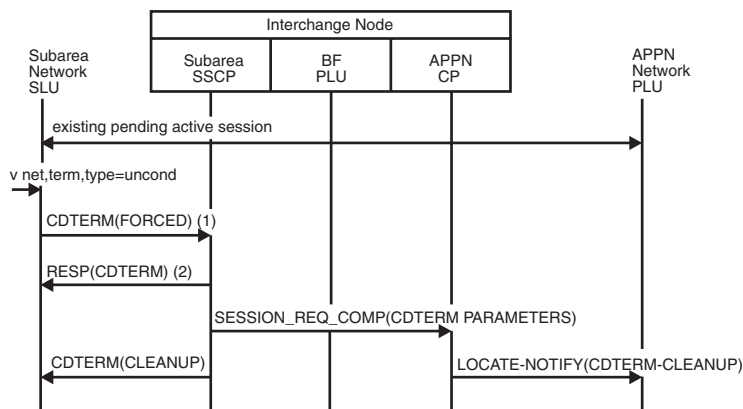


Figure 218. SA(SLU)==ICN...APPN network(PLU), forced termination of pending active session

1. The CDTERM type depends on the V NET,TERM type.

V NET,TERM type
CDTERM type

COND
ORDERLY

UNCOND
FORCED

FORCE
CLEANUP

2. APPN has only orderly and cleanup termination. Therefore, the forced termination is promoted to clean up when it crosses from subarea into APPN.

SA(PLU)==ICN...APPN network(SLU), forced termination of pending active session (PLU accessible without going into APPN)

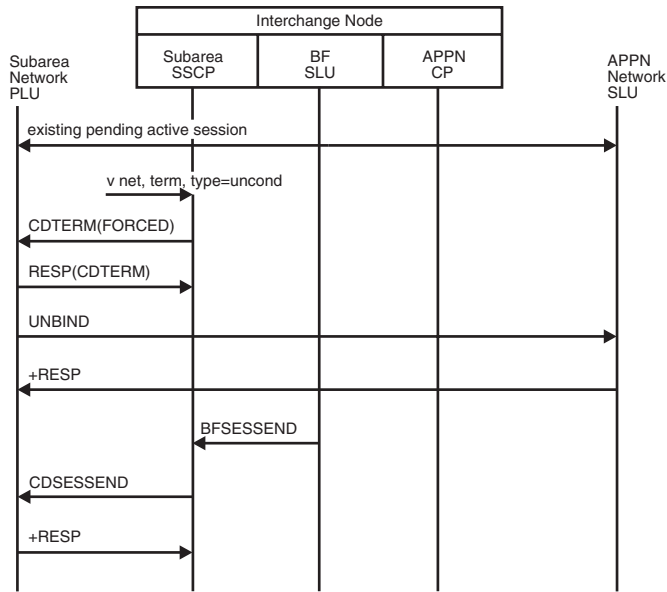


Figure 219. SA(PLU)==ICN...APPN network(SLU), forced termination of pending active session (PLU accessible without going into APPN)

Note: Whenever a forced termination crosses the boundary from subarea into APPN, it is promoted to clean up. In this case because the PLU is accessible without going into the APPN network, promotion does not occur.

SA(PLU)==ICN...APPN network(SLU), forced termination of queued session

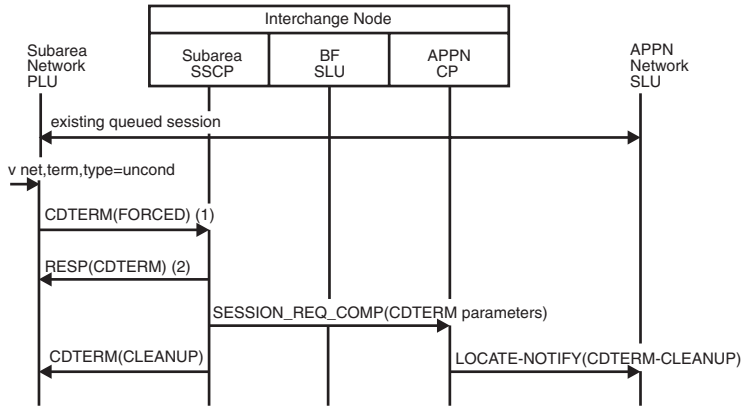


Figure 220. SA(PLU)==ICN...APPN network(SLU), forced termination of queued session

1. The CDTERM type depends on the V NET,TERM type.

V NET,TERM type
CDTERM type

COND
ORDERLY

UNCOND
FORCED

FORCE
CLEANUP

2. APPN has only orderly and cleanup termination. Therefore, the forced termination is promoted to cleanup when it crosses from subarea into APPN.

SA(PLU)==ICN...APPN network(SLU), forced termination of queued session (PLU accessible without going into APPN)

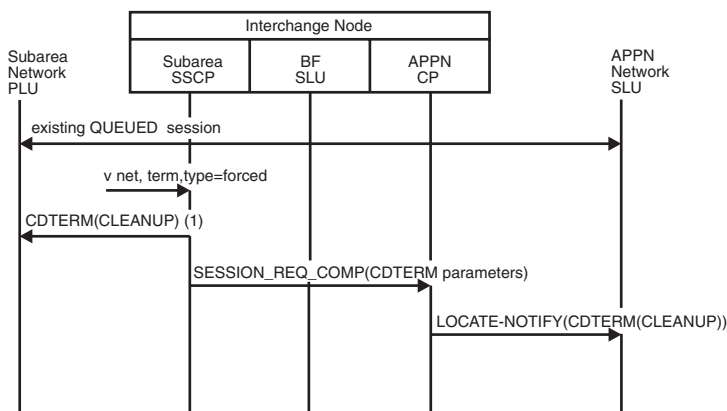


Figure 221. SA(PLU)==ICN...APPN network(SLU), forced termination of queued session (PLU accessible without going into APPN)

1. Because the session is queued (instead of pending active) and the forced termination is not issued in the primary LU domain, a CLEANUP is sent.

SA (PLU)==ICN...APPN network (SLU), session release request

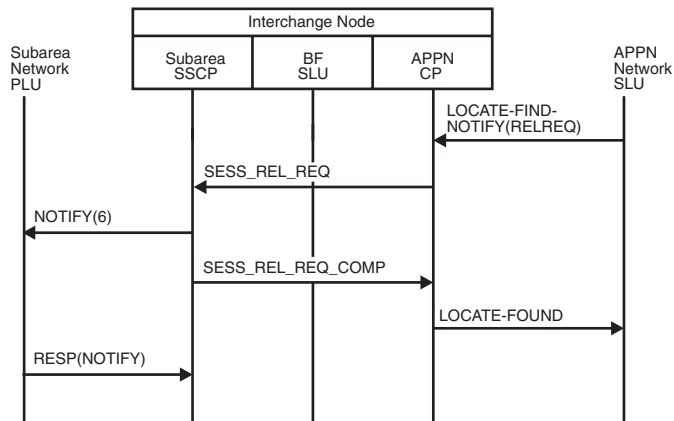


Figure 222. SA (PLU)==ICN...APPN network (SLU), session release request

SA (SLU)==ICN...APPN network (PLU), session release request

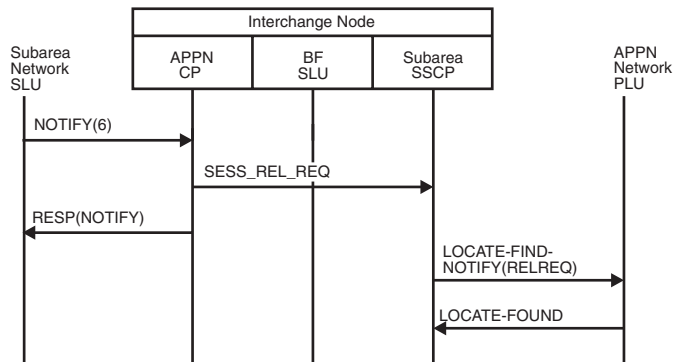


Figure 223. SA (SLU)==ICN...APPN network (PLU), session release request

CLSDST PASS through APPN

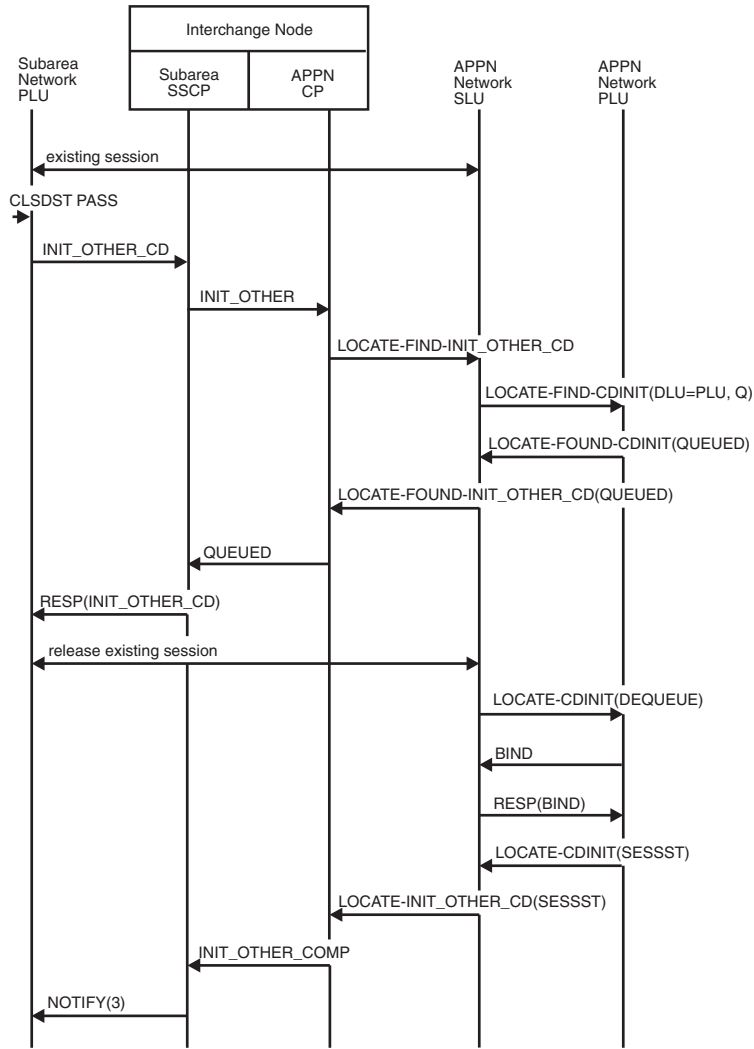


Figure 224. CLSDST PASS through APPN. The SLU is single-session capable.

CLSDST PASS from APPN to subarea

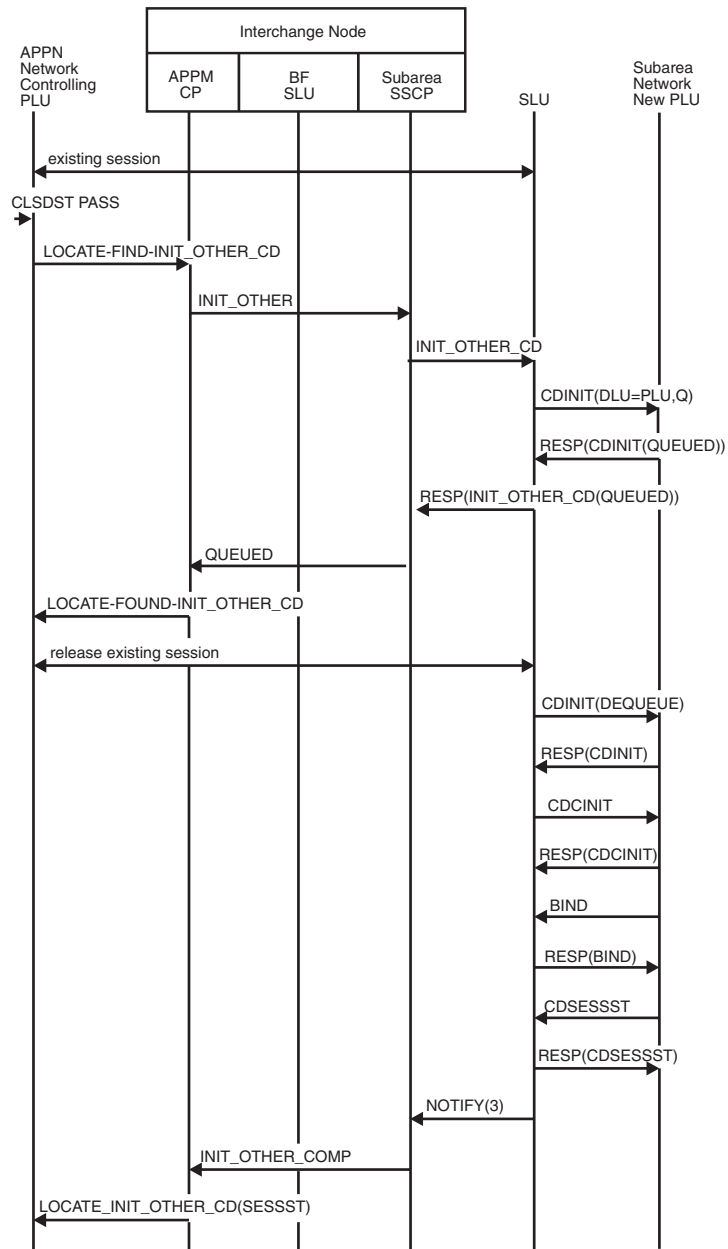


Figure 225. CLSDST PASS from APPN to subarea. The SLU is single-session capable.

EN-NN-EN, PLU-initiated, no queuing

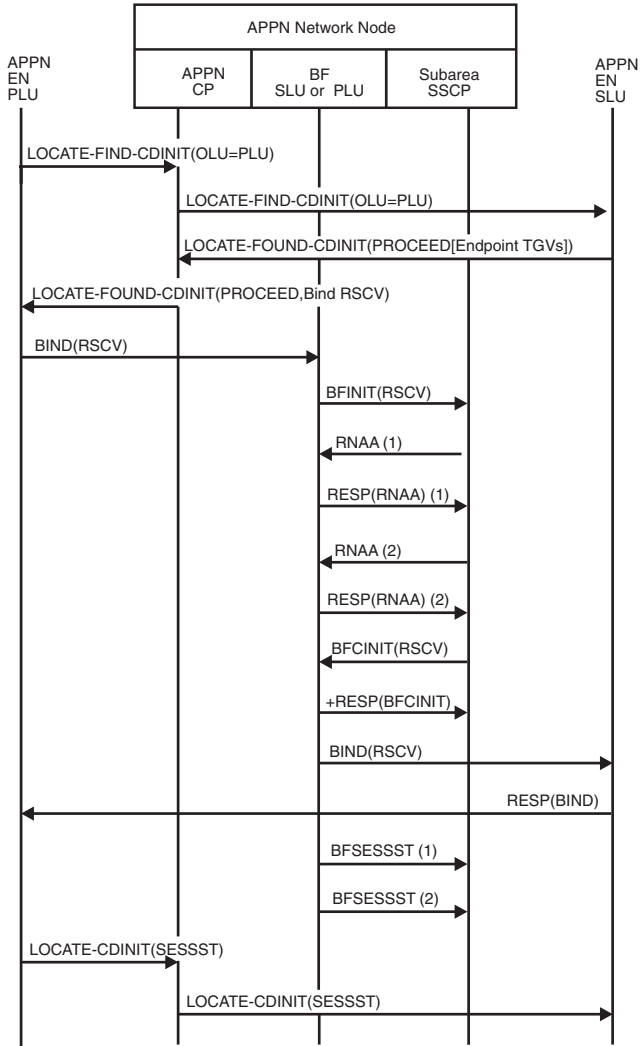


Figure 226. EN-NN-EN, PLU-initiated, no queuing (Including BIND flows for intermediate network node)

1. For the PLU side of the session
2. For the SLU side of the session

Intermediate network node (INN) BIND

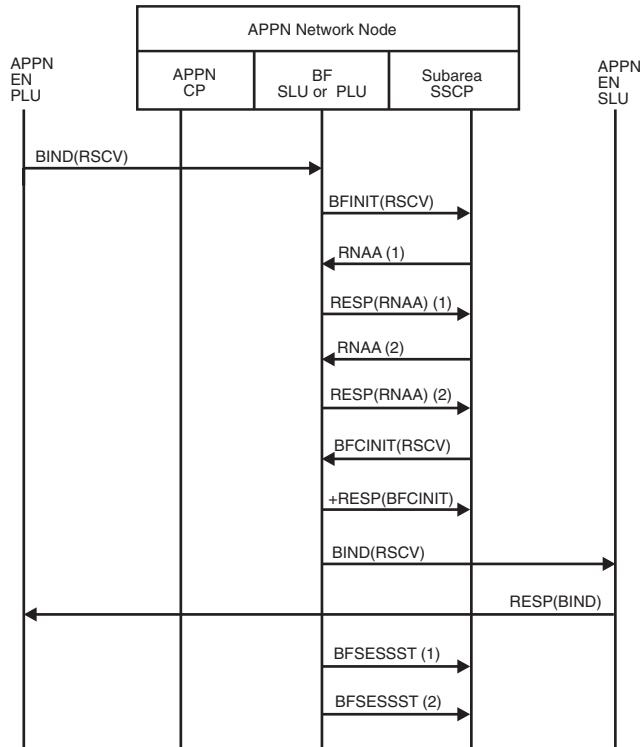


Figure 227. Intermediate network node (INN) BIND. The LOCATE did not go through this node.

1. For the PLU side of the session
2. For the SLU side of the session

APPN network (PLU)...ICN==SA(SLU), PLU-initiated, search-only flow transformed into a DSRLST

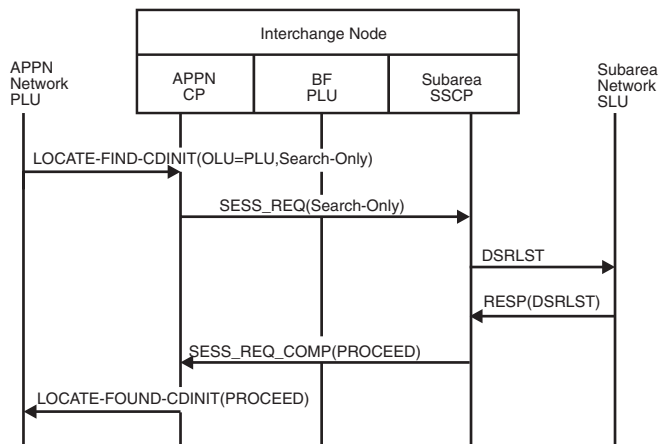


Figure 228. APPN network (PLU)...ICN==SA(SLU), PLU-initiated, search-only flow transformed into a DSRLST

APPN network (PLU)...ICN==SA(SLU), PLU-initiated, no queuing

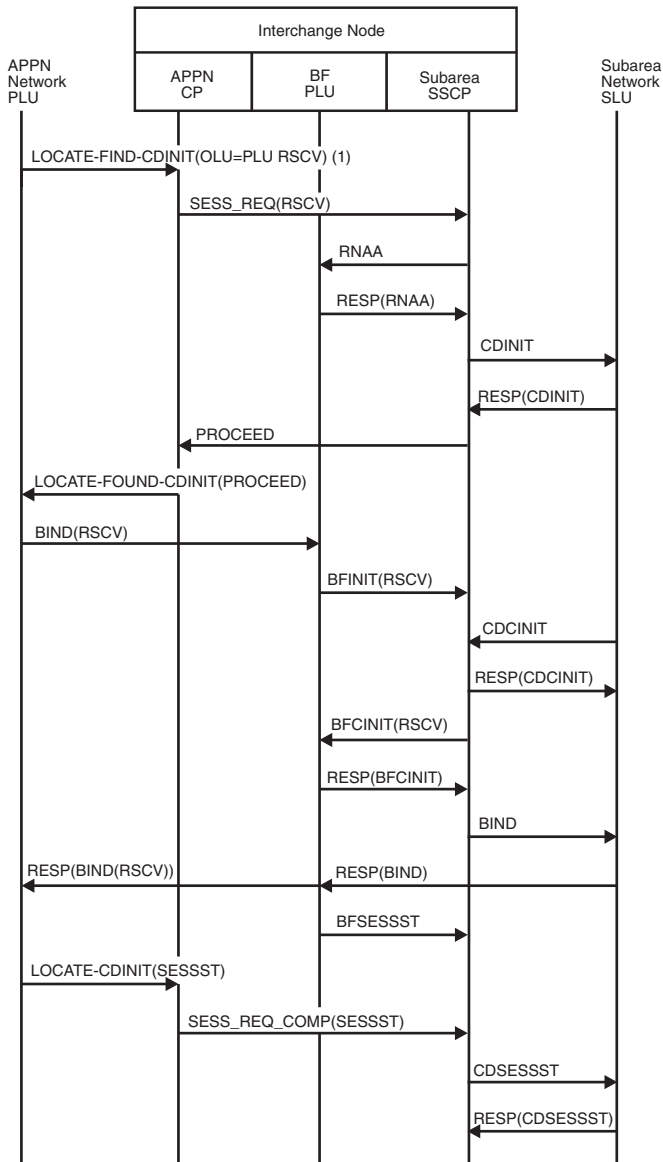


Figure 229. APPN network (PLU)...ICN==SA(SLU), PLU-initiated, no queueing

1. Because the DLU is in subarea, the NNS(OLU) precomputed the RSCV.

APPN network (PLU)...ICN==SA(SLU), PLU-initiated, directed search without required precomputed RSCV

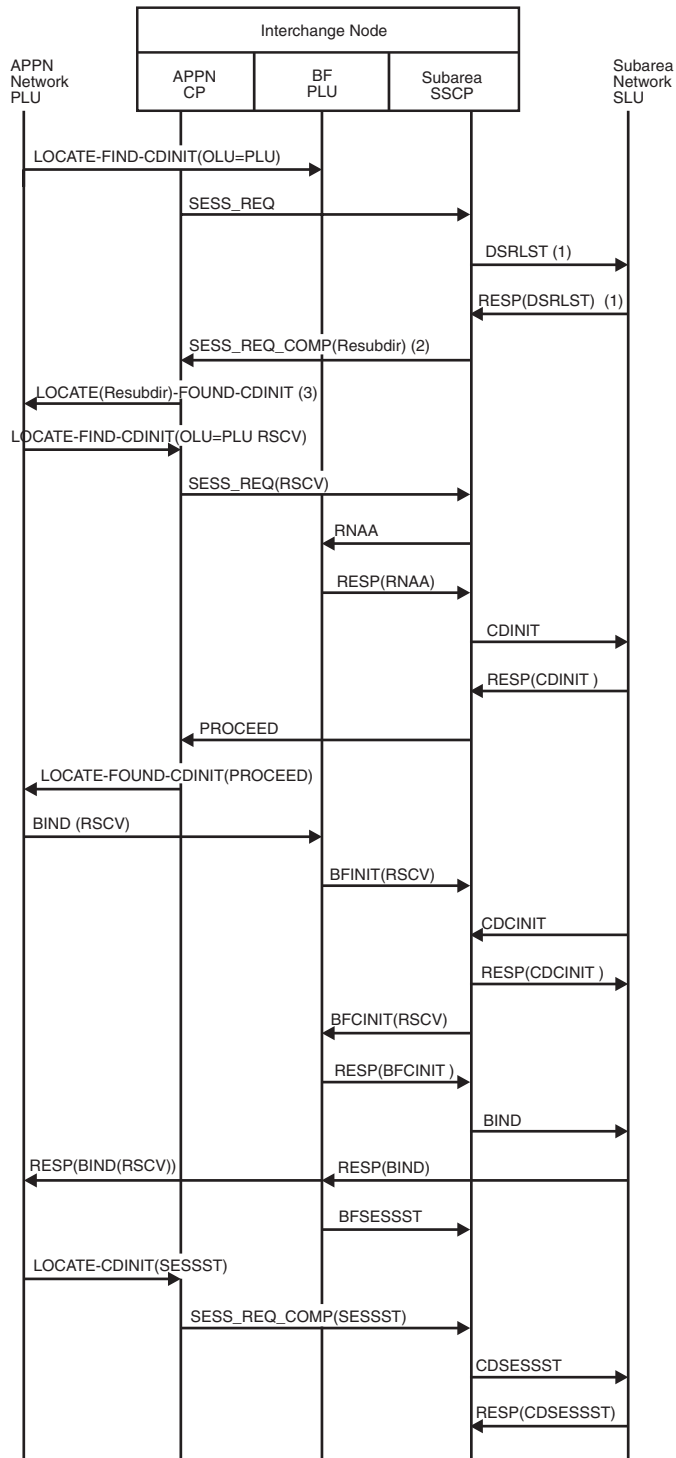


Figure 230. APPN network (PLU)...ICN==SA(SLU), PLU-initiated, directed search without required precomputed RSCV

1. Optional; sent if owning CP is not known.
2. Resubdir means "Resubmit Request on a Directed Search."
3. The interchange node returns the fact that the SLU is in a subarea network and requires a precomputed RSCV.

APPN network (PLU)...ICN==SA(SLU), PLU-initiated, USERVAR resolution required

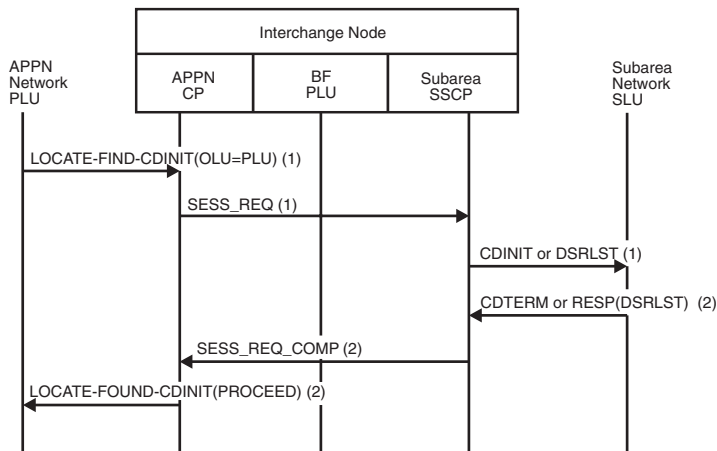


Figure 231. APPN network (PLU)...ICN==SA(SLU), PLU-initiated, USERVAR resolution required

1. Generic USERVAR name
2. Resolved USERVAR name

For remaining session setup flows, see Figure 229 on page 579.

APPN network (PLU)...ICN==SA(SLU), PLU-initiated, queued by the SLU

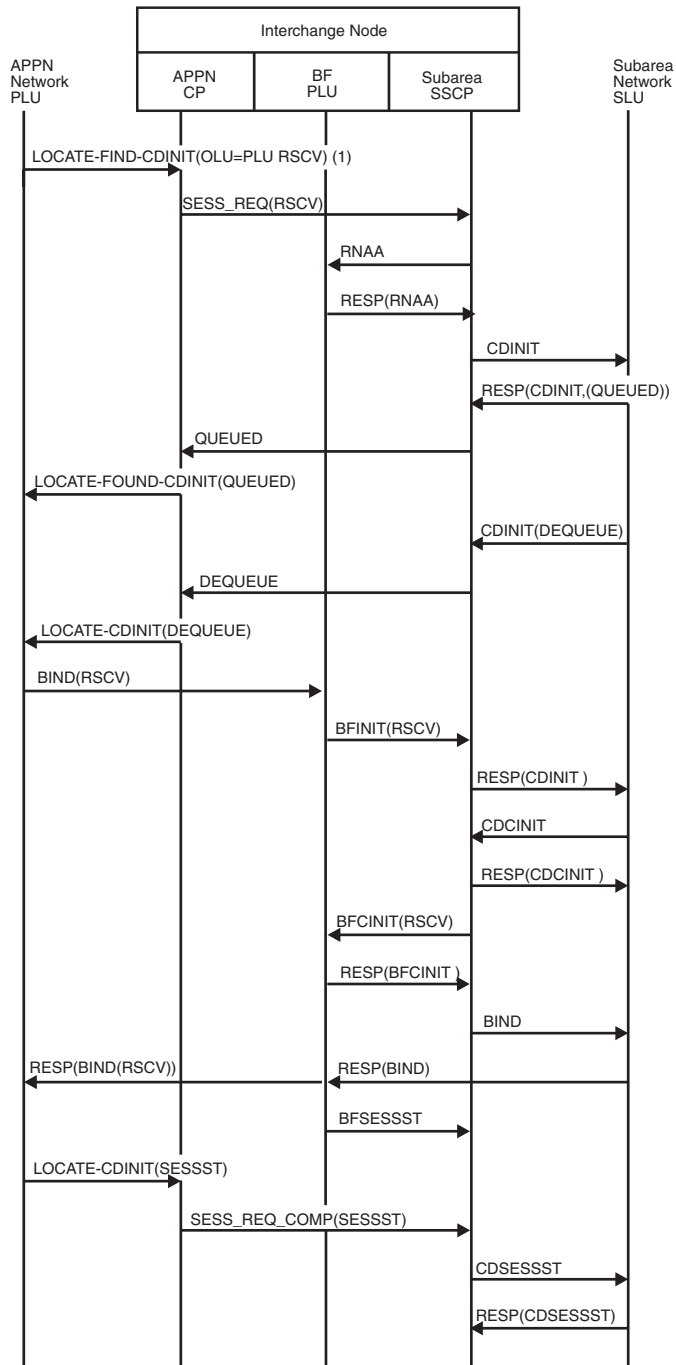


Figure 232. APPN network (PLU)...ICN==SA(SLU), PLU-initiated, queued by the SLU

1. Because the DLU is in subarea, the NNS(OLU) precomputed the RSCV.

APPN network (SLU)...ICN==SA(PLU), SLU-initiated, no queueing

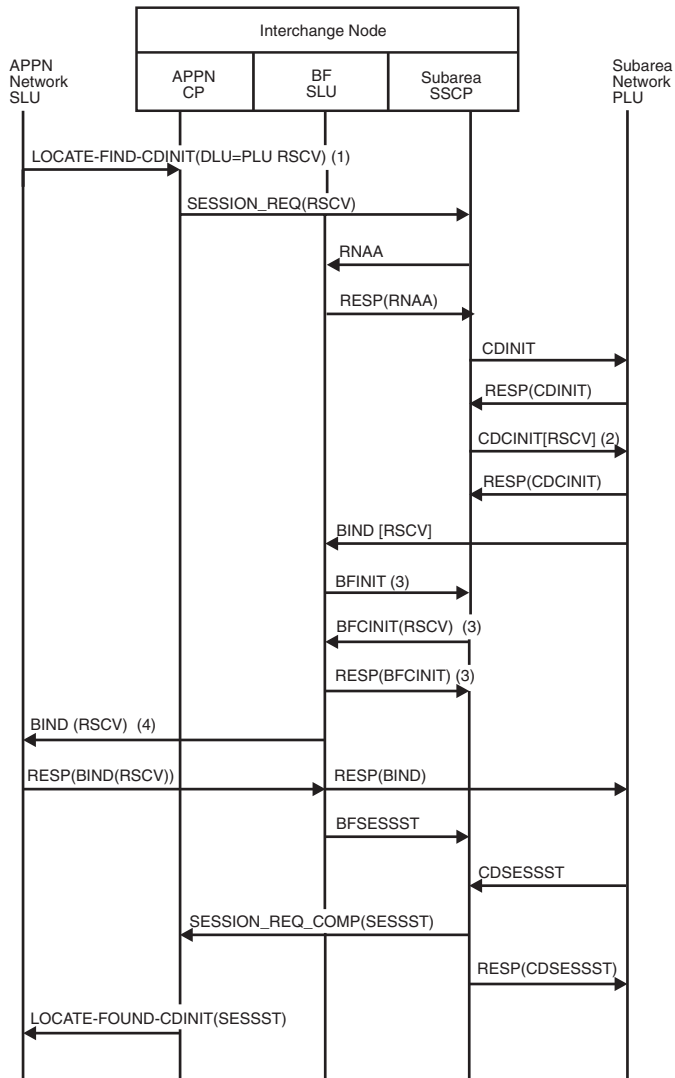


Figure 233. APPN network (SLU)...ICN==SA(PLU), SLU-initiated, no queueing

1. Because the DLU is in subarea, the NNS(OLU) precomputes the RSCV.
2. If the adjacent SSCP toward the PLU is VTAM V4R1 or higher and has the same network identifier, the RSCV is passed on the CDCINIT.
3. If the RSCV is passed on the CDCINIT, these flows will not occur.
4. The BIND does not have to follow the same path as the LOCATE flows.

APPN network (SLU)...ICN==SA(PLU), SLU-initiated, queued by the PLU

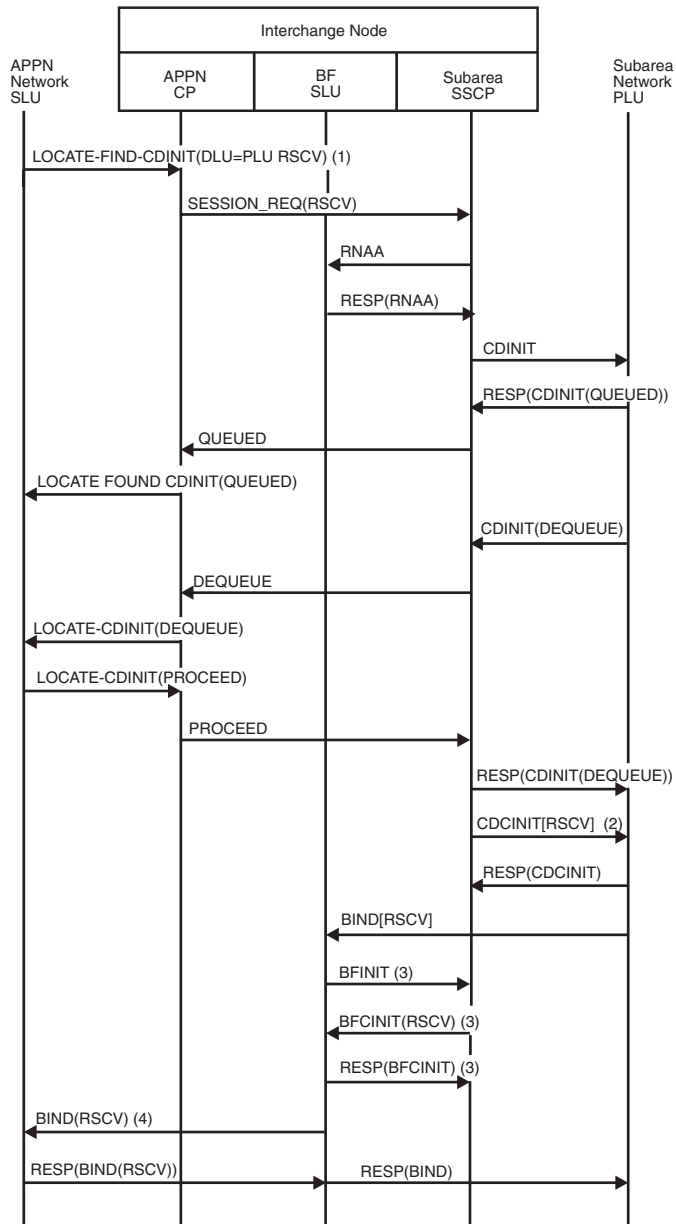


Figure 234. APPN network (SLU)...ICN==SA(PLU), SLU-initiated, queued by the PLU (part 1 of 2)

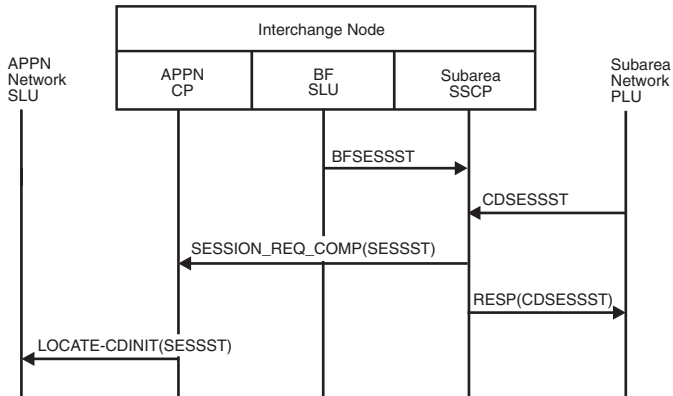


Figure 235. APPN network (SLU)...ICN==SA(PLU), SLU-initiated, queued by the PLU (part 2 of 2)

1. Because the DLU is in subarea, the NNS(OLU) precomputed the RSCV.
2. If the adjacent SSCP into the subarea is VTAM V4R1 or higher and has the same network identifier, the RSCV is passed on the CDCINIT.
3. If the RSCV is passed on the CDCINIT, these flows will not occur.
4. The BIND does not have to follow the same path as the LOCATE flows.

APPN network (SLU)...ICN==SA(PLU), autologon (PLU not available initially)

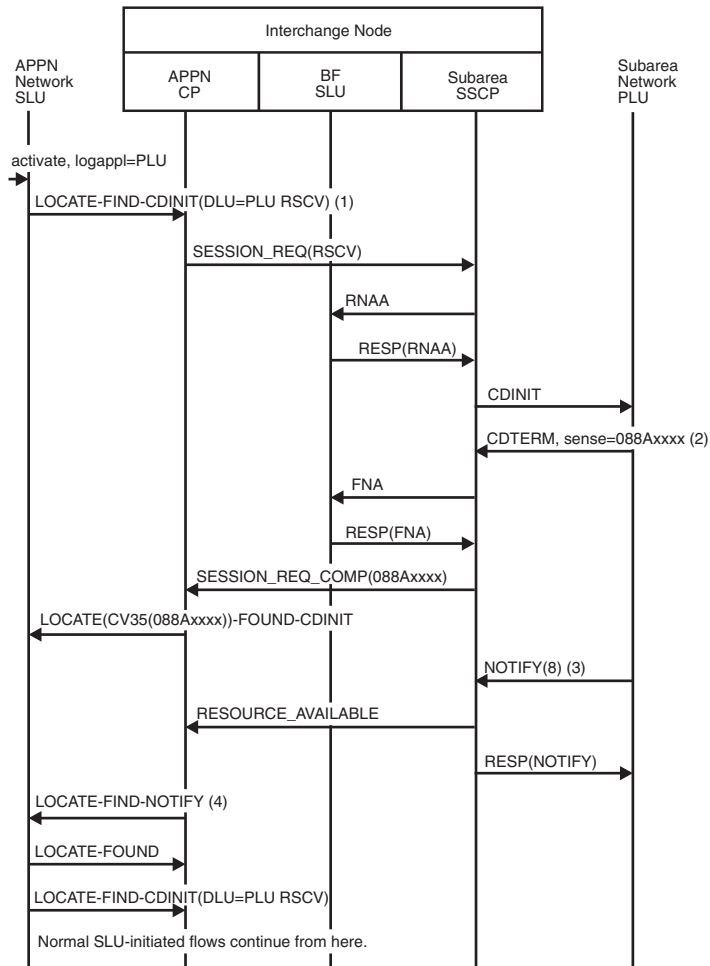


Figure 236. APPN network (SLU)...ICN==SA(PLU), autologon (PLU not available initially)

1. Because the DLU is in subarea, the NNS(OLU) precomputed the session RSCV.
2. The PLU is not currently available.
3. Some time later, the PLU becomes available.
4. Resource enabled.

APPN network (PLU)...ICN==VR-based TG==ICN...APPN network (SLU)

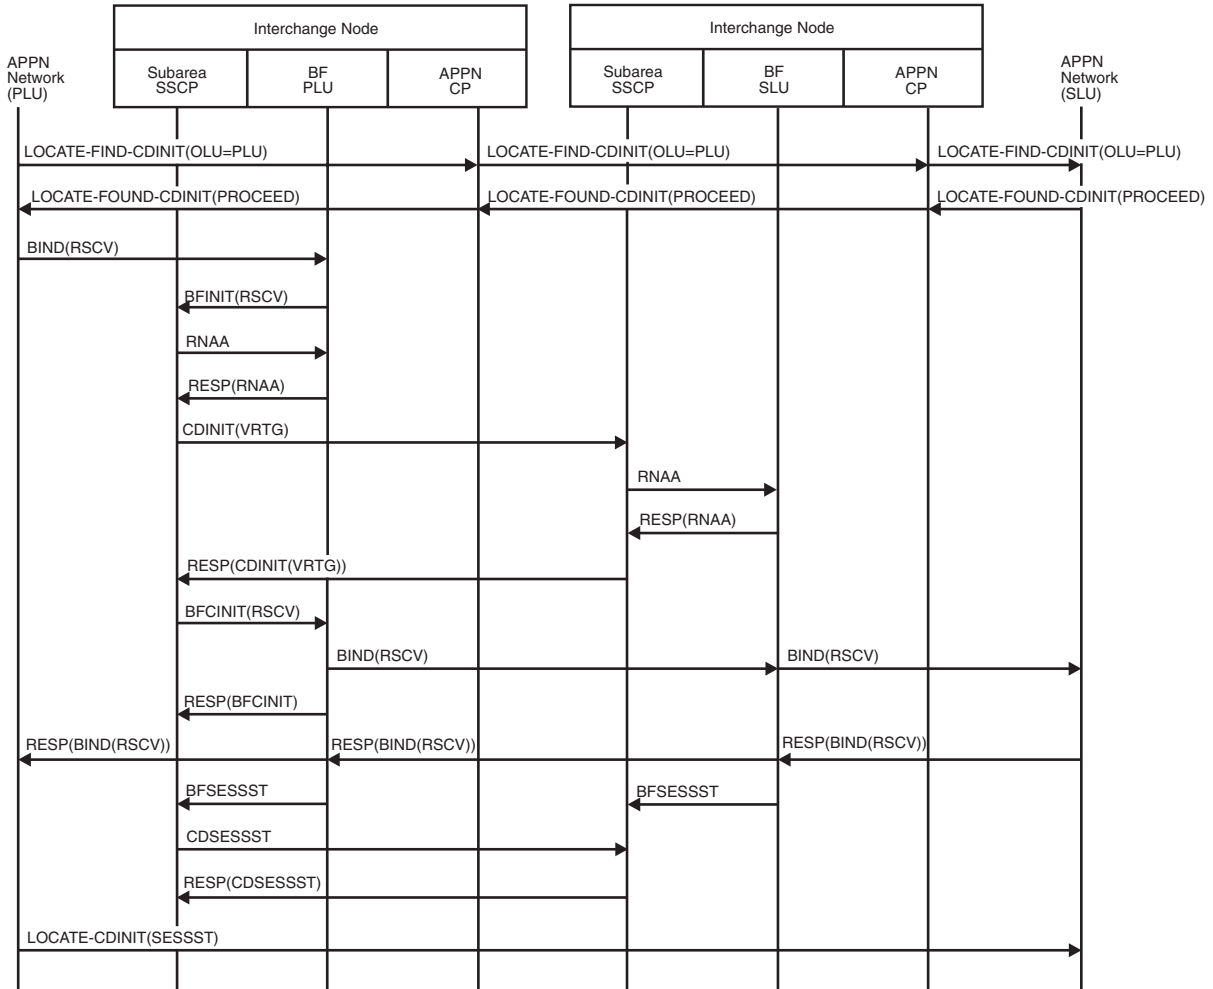


Figure 237. APPN network (PLU)...ICN==VR-based TG==ICN...APPN network (SLU), PLU-initiated

Dependent LU server flows

Figure 238 on page 589 through Figure 260 on page 611 illustrate the flow of requests and responses between dependent logical unit requestors and servers.

Index of dependent LU server flows

Table 43 lists the dependent LU server flows illustrated here.

Table 43. Index of dependent LU server flows

Single subnetwork flow	Page
Activating resources	
CPSVRMGR pipe activation, DLUR-Initiated	Figure 238 on page 589
CPSVRMGR pipe activation, DLUS-initiated	Figure 239 on page 590
Dependent LUs, dynamic registration and activation of	Figure 241 on page 592
Dependent LUs, activation of predefined	Figure 242 on page 592

Table 43. Index of dependent LU server flows (continued)

Single subnetwork flow	Page
Figure 240 on page 591	
SSCP-PU session activation race	Figure 243 on page 593
Deactivating resources	
CPSVRMGR pipe deactivation	Figure 244 on page 594
Downstream PU outage	Figure 245 on page 595
REQDISCONT (immediate) received from downstream PU	Figure 247 on page 597
REQDISCONT (normal) received from downstream PU	Figure 246 on page 596
LU-LU sessions	
APPN PLU-initiated to a dependent SLU	Figure 252 on page 602
Session termination, USS flows for	Figure 255 on page 605
USS SLU-initiated to APPN PLU	Figure 253 on page 603
USS SLU-initiated to subarea PLU	Figure 254 on page 604
SSCP-PU, SSCP-LU session deactivation	
Forced	Figure 249 on page 599
Normal	Figure 248 on page 598
With Giveback (ANS=CONT)	Figure 251 on page 601
With Giveback (ANS=STOP)	Figure 250 on page 600
Cross Subnetwork Flow	
Page	
PLU-Initiated Session with DLUS and DLUR within Different Subnetworks	Figure 256 on page 606
PLU-Initiated Session with DLUS and PLU in one Subnetwork and DLUR in Another	Figure 258 on page 609
SLU-Initiated Session with DLUS and DLUR within Different Subnetworks	Figure 260 on page 611

Single subnetwork flows

Figure 238 on page 589 through Figure 255 on page 605 show the flow of requests and responses between dependent logical unit requestors and servers within a single subnetwork.

DLUR-initiated CPSVRMGR pipe activation

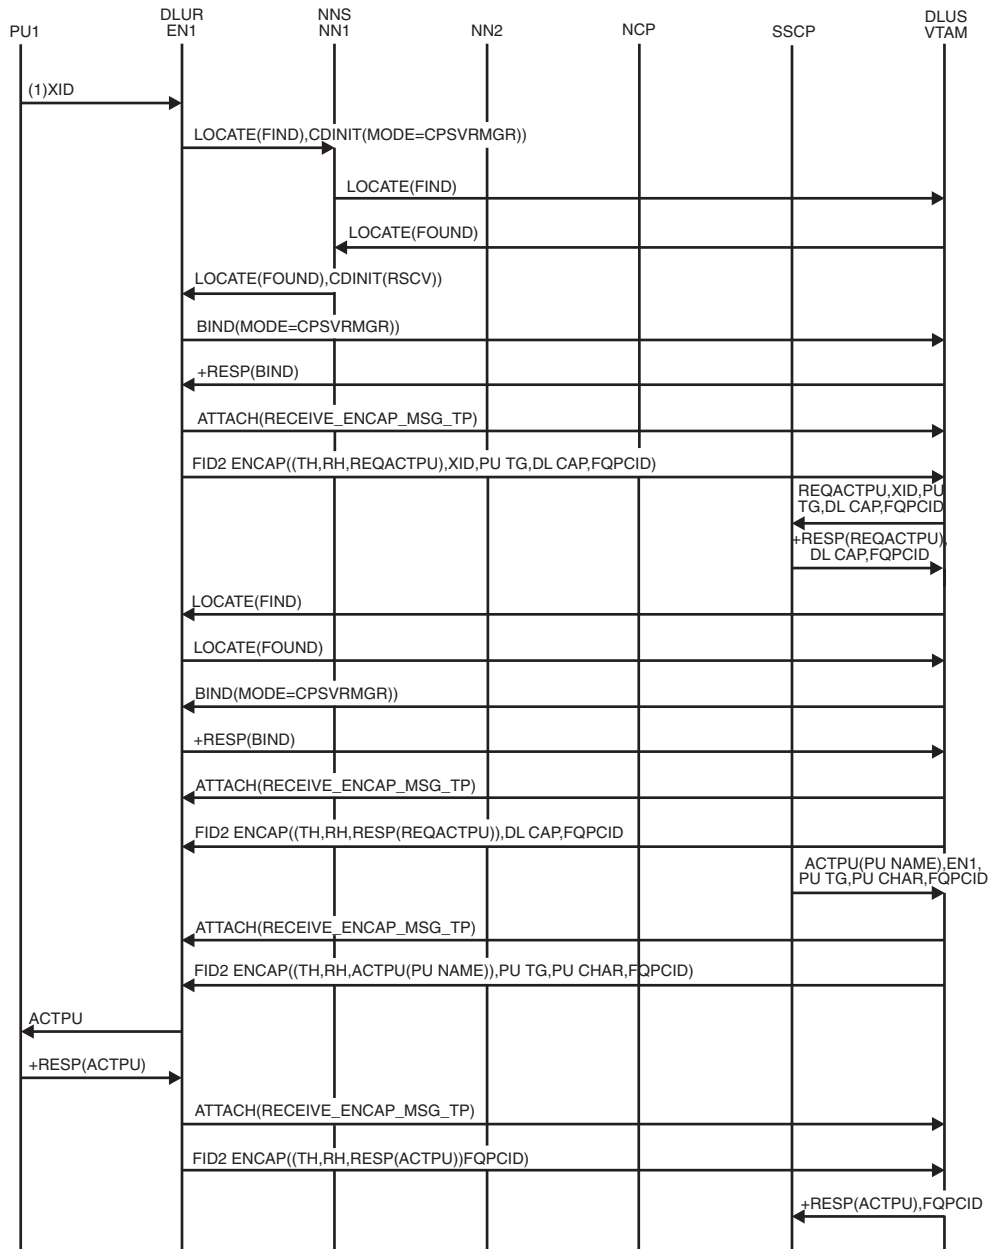


Figure 238. DLUR-initiated CPSVRMGR pipe activation

1. XID flows as a result of a set normal response mode (SNRM) RU, an external command, or an internal activation signal.

DLUS-initiated CPSVRMGR pipe activation

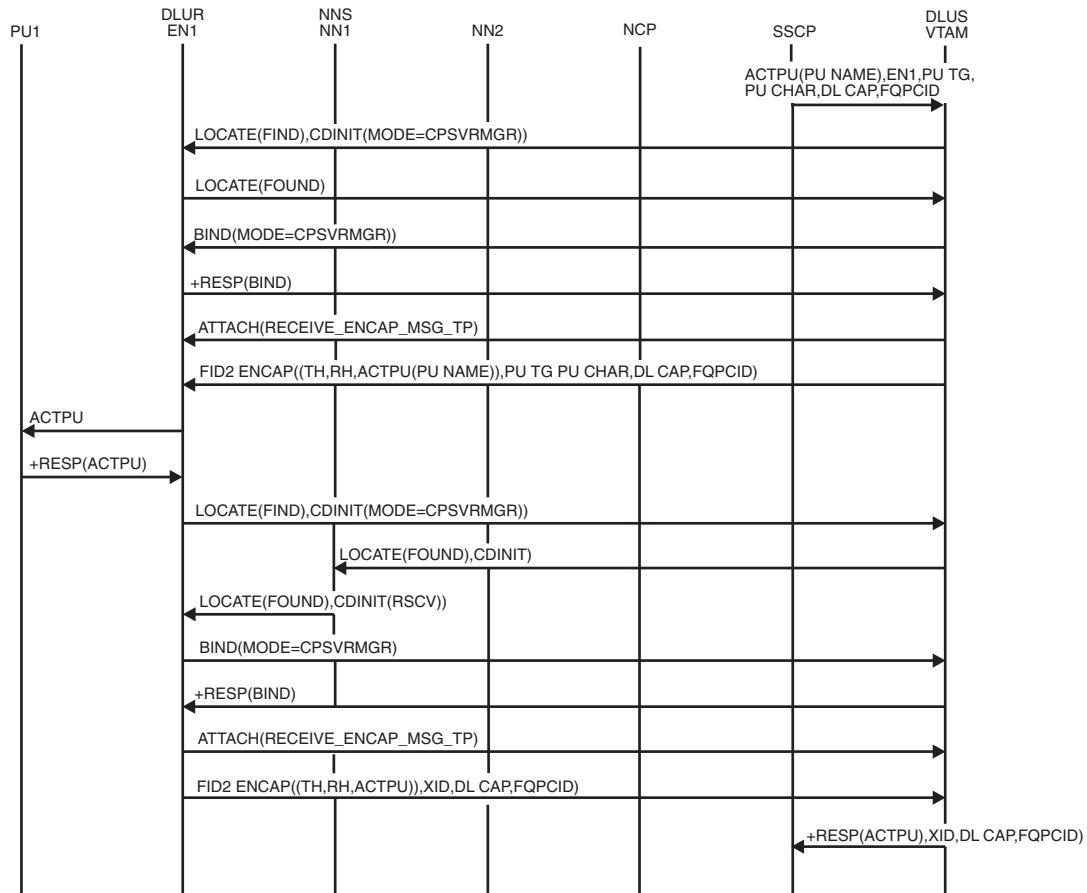


Figure 239. DLUS-initiated CPSVRMGR pipe activation

Dynamic PU activation

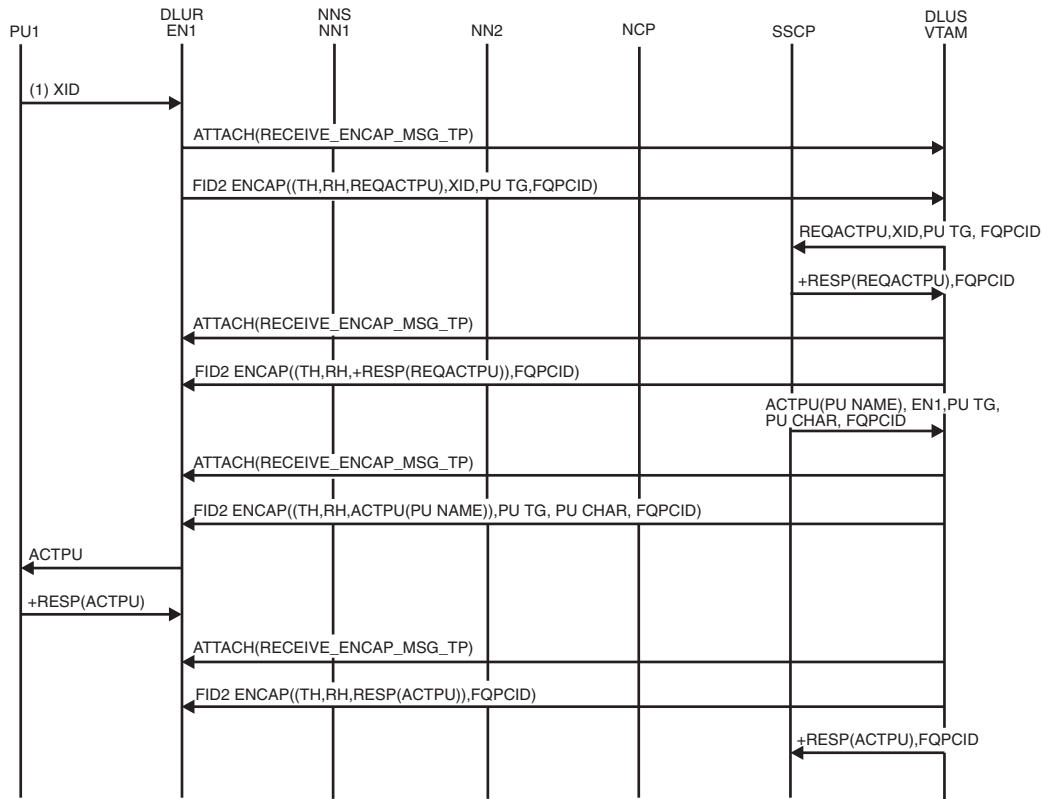


Figure 240. Dynamic PU activation

1. XID flows as a result of a set normal response mode (SNRM) RU, an external command, or an internal activation signal.

Dynamic registration and activation of dependent LUs

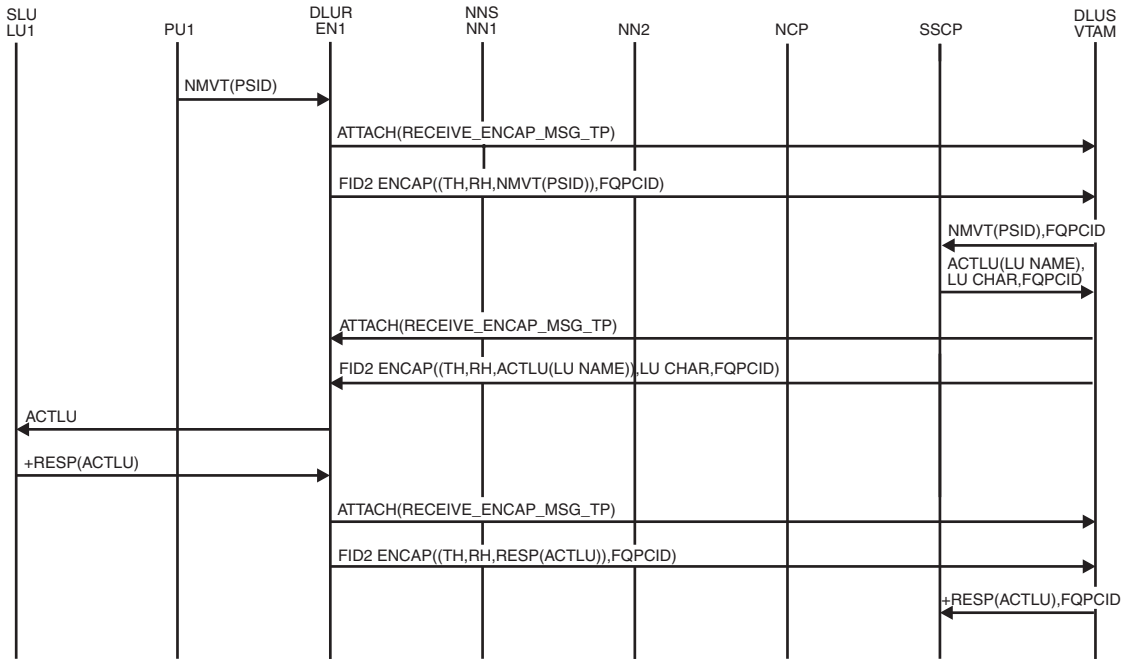


Figure 241. Dynamic registration and activation of dependent LUs

Activation of predefined dependent LUs

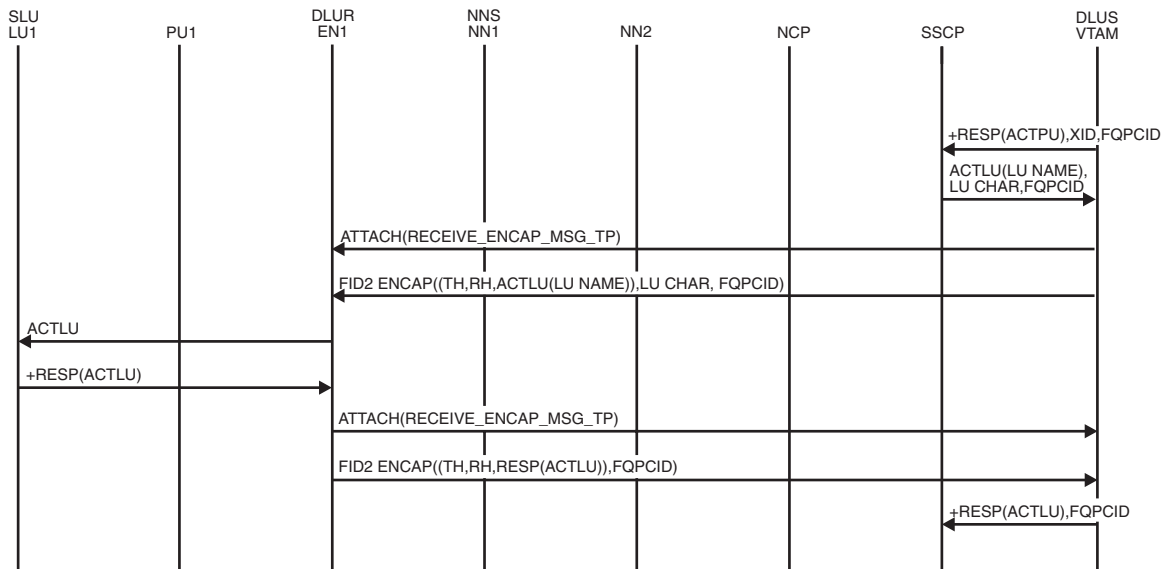


Figure 242. Activation of predefined dependent LUs

SSCP-PU session activation race

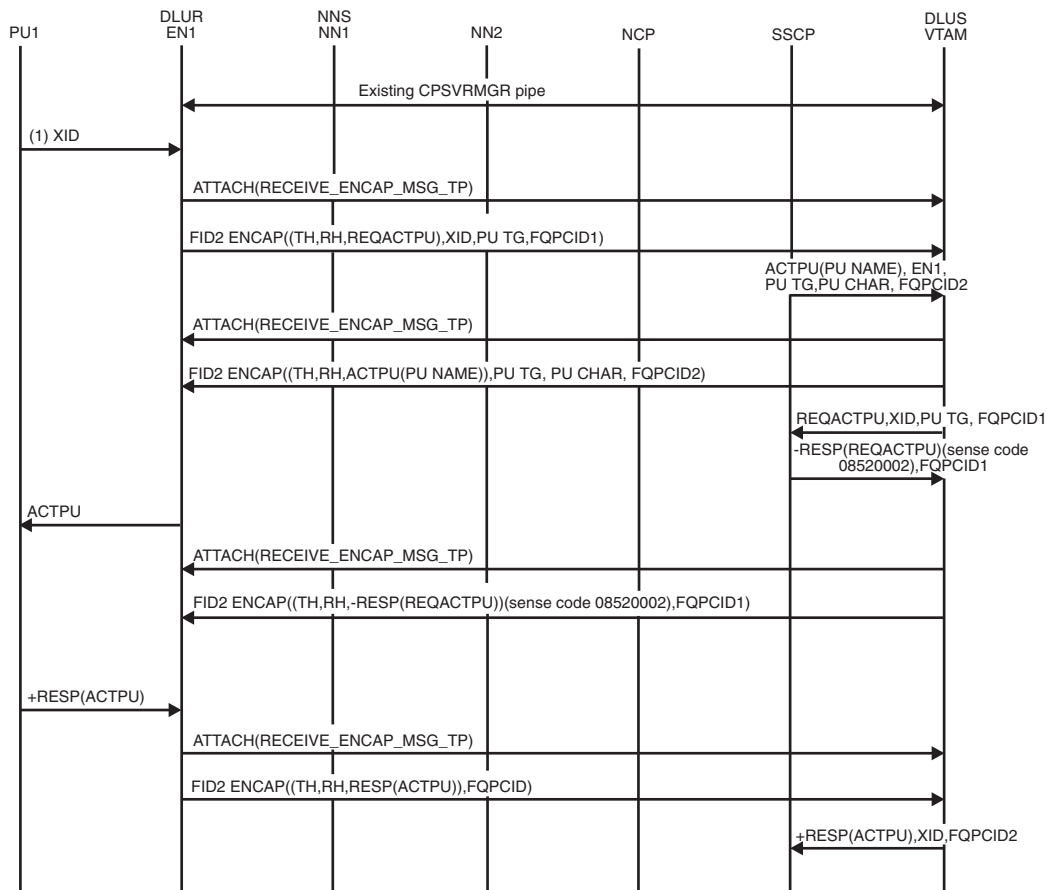


Figure 243. SSCP-PU session activation race

1. XID flows as a result of a set normal response mode (SNRM) RU, an external command, or an internal activation signal.

CPSVRMGR pipe deactivation

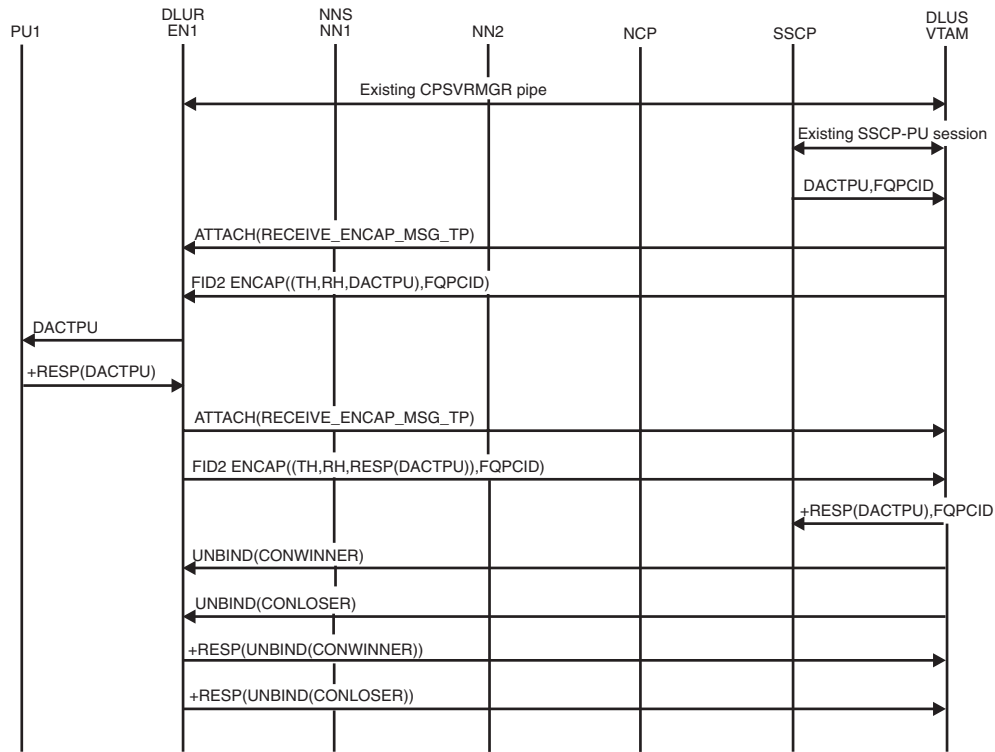


Figure 244. CPSVRMGR pipe deactivation

Downstream PU outage

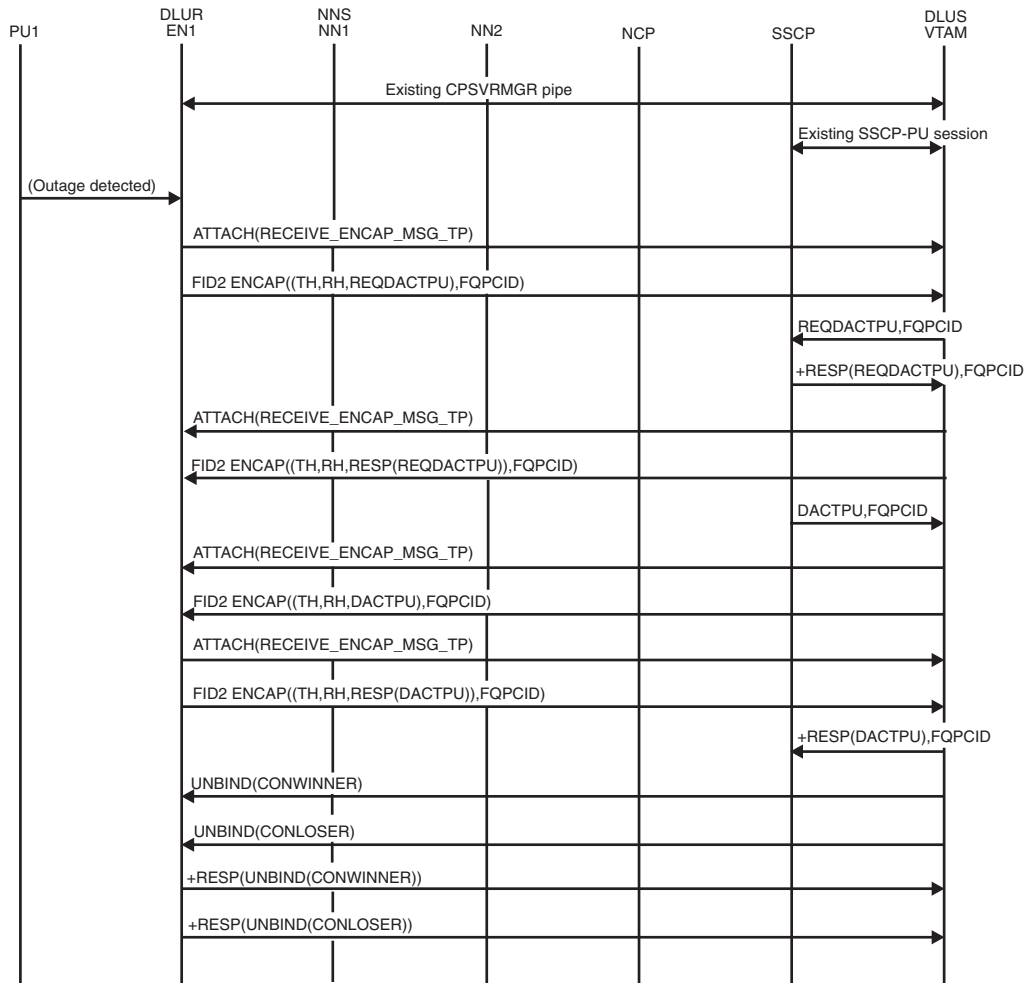


Figure 245. Downstream PU outage

Receipt of REQDISCONT (normal) from downstream PU

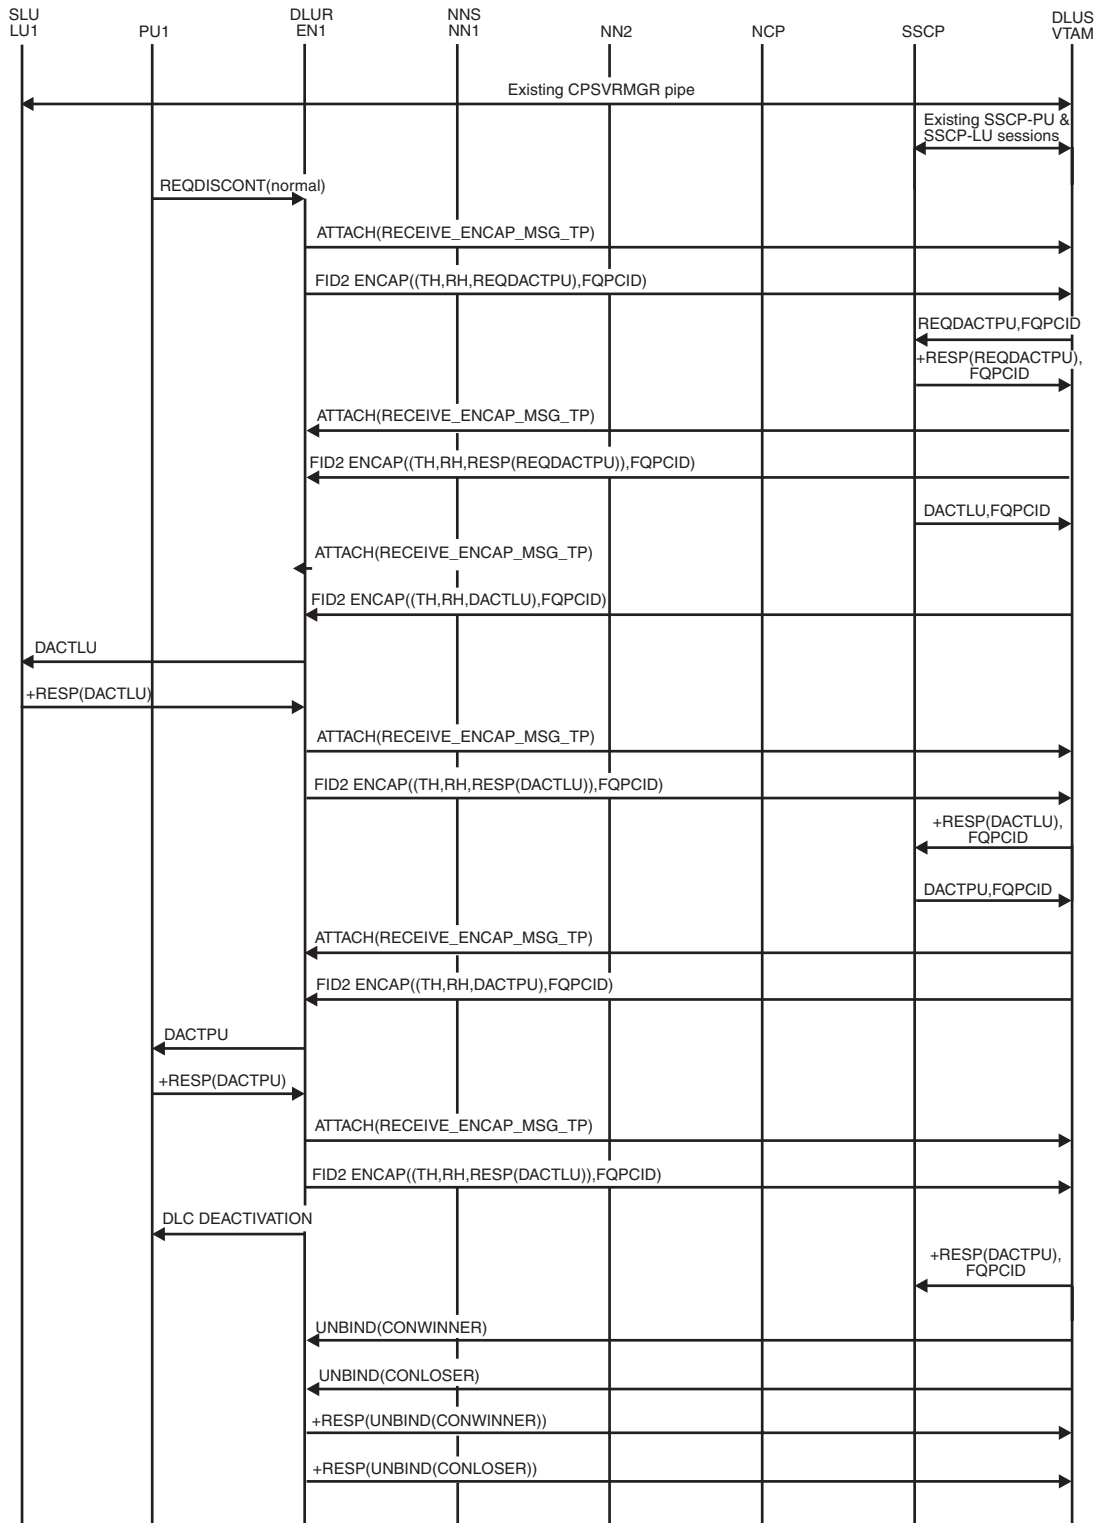


Figure 246. Receipt of REQDISCONT (normal) from downstream PU

Receipt of REQDISCONT (immediate) from downstream PU

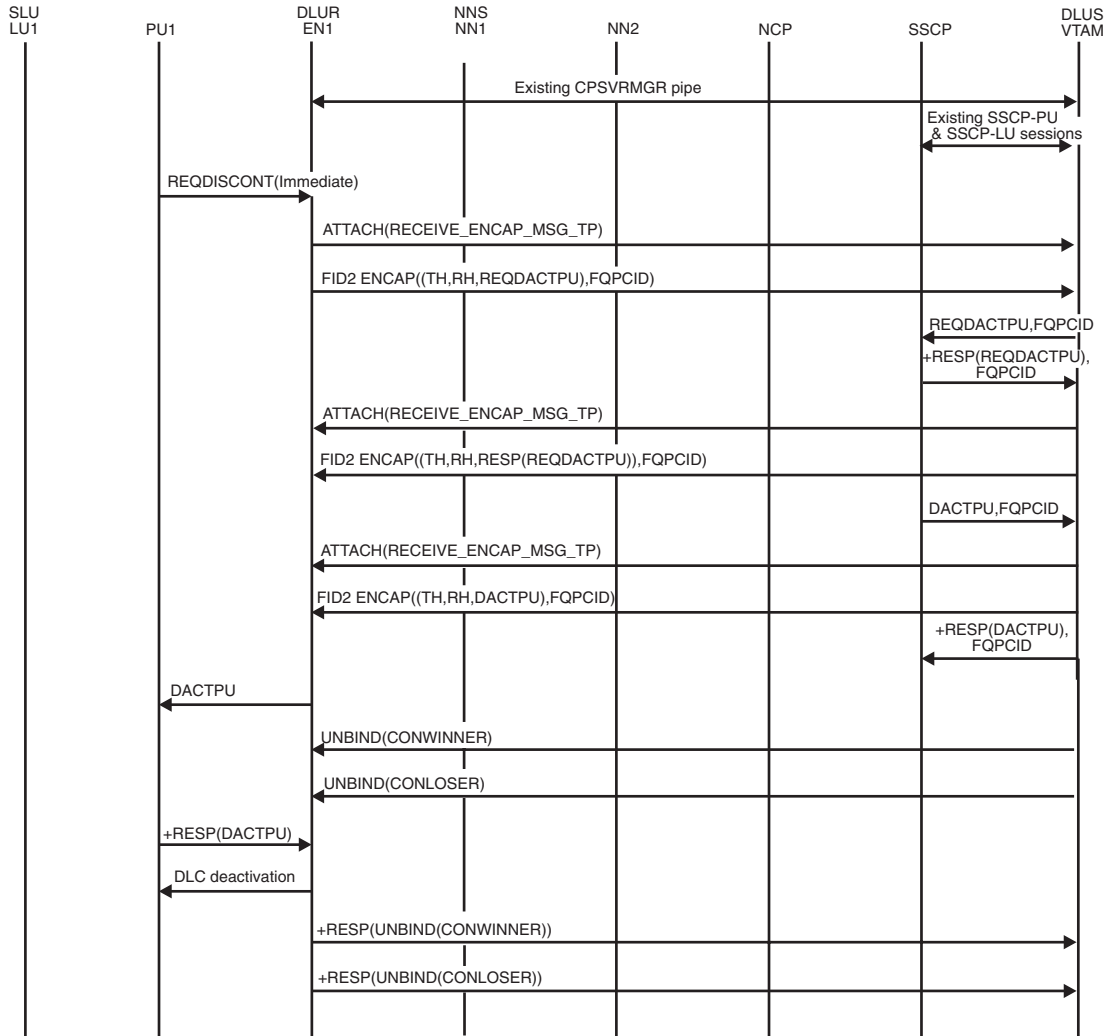


Figure 247. Receipt of REQDISCONT (immediate) from downstream PU

Normal SSCP-PU/SSCP-LU session deactivation

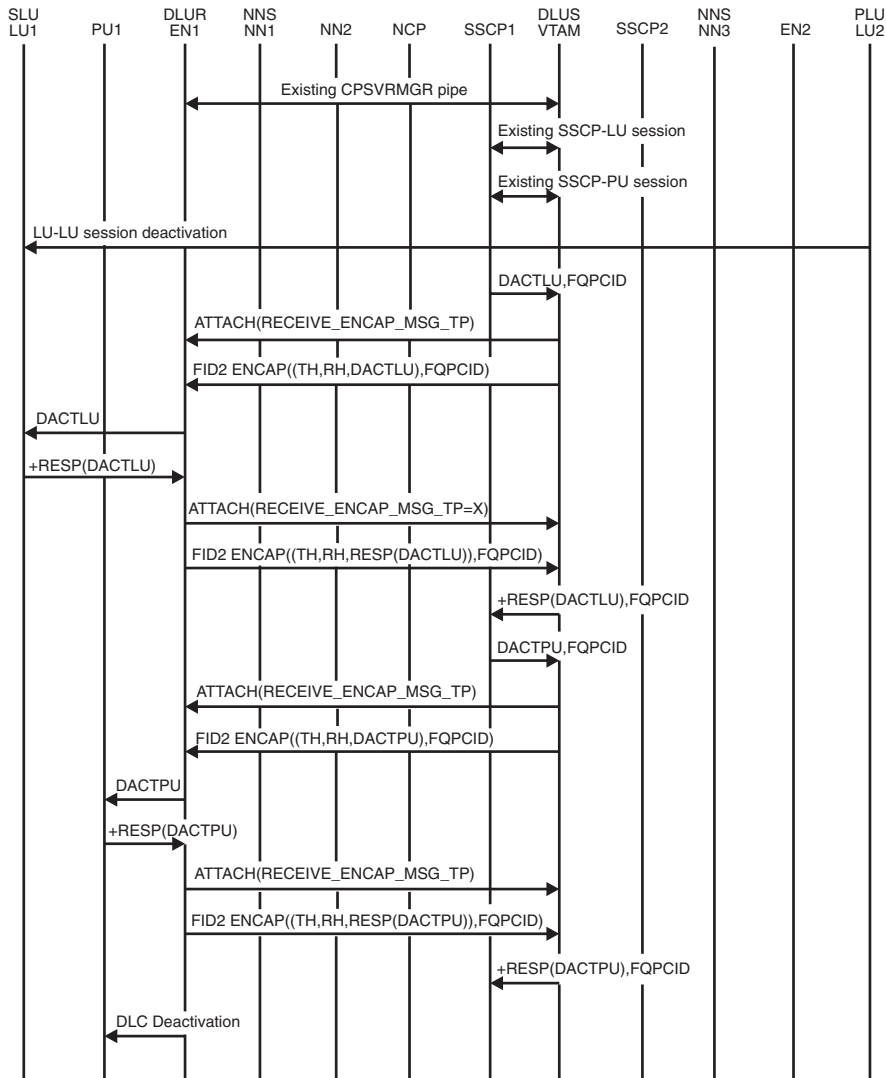


Figure 248. Normal SSCP-PU/SSCP-LU session deactivation

Forced SSCP-PU/SSCP-LU session deactivation

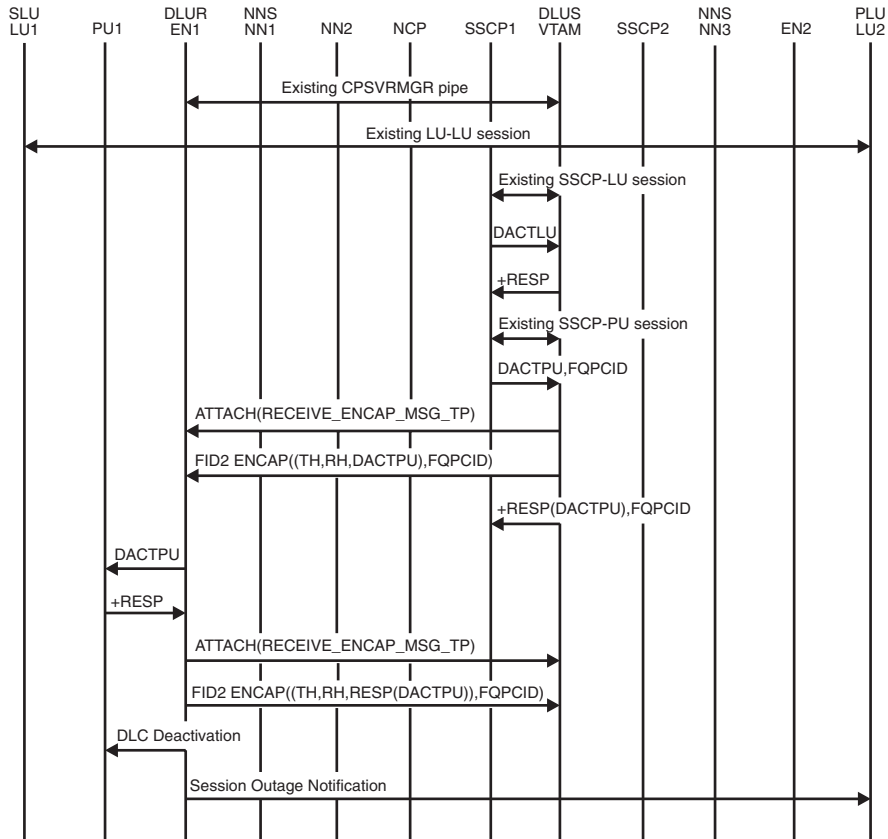


Figure 249. Forced SSCP-PU/SSCP-LU session deactivation

Giveback SSCP-PU/SSCP-LU session deactivation (ANS=STOP)

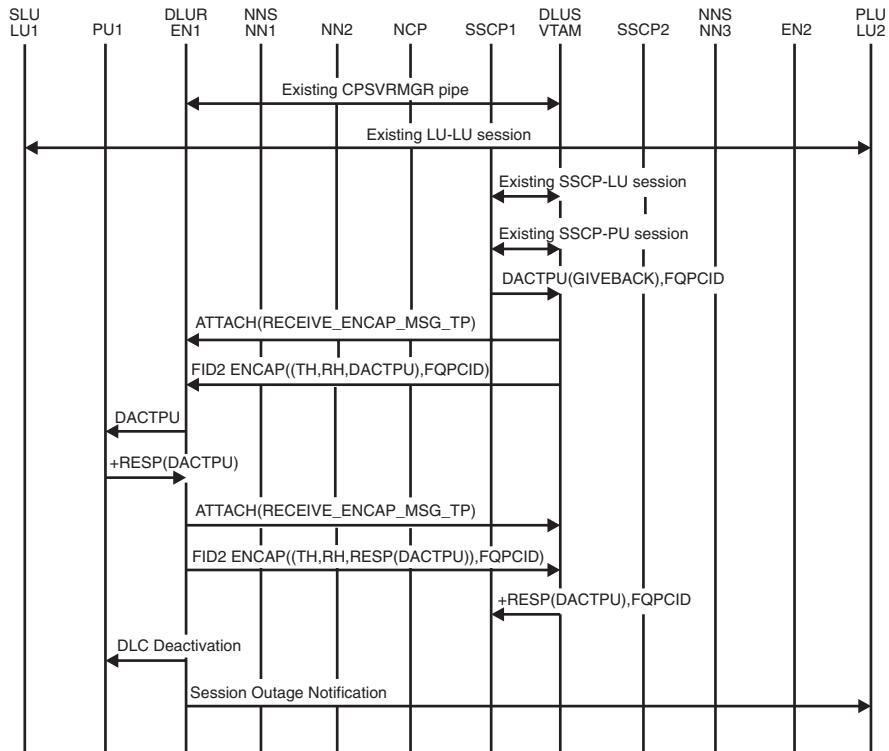


Figure 250. Giveback SSCP-PU/SSCP-LU session deactivation (ANS=STOP)

Giveback SSCP-PU/SSCP-LU session deactivation (ANS=CONT)

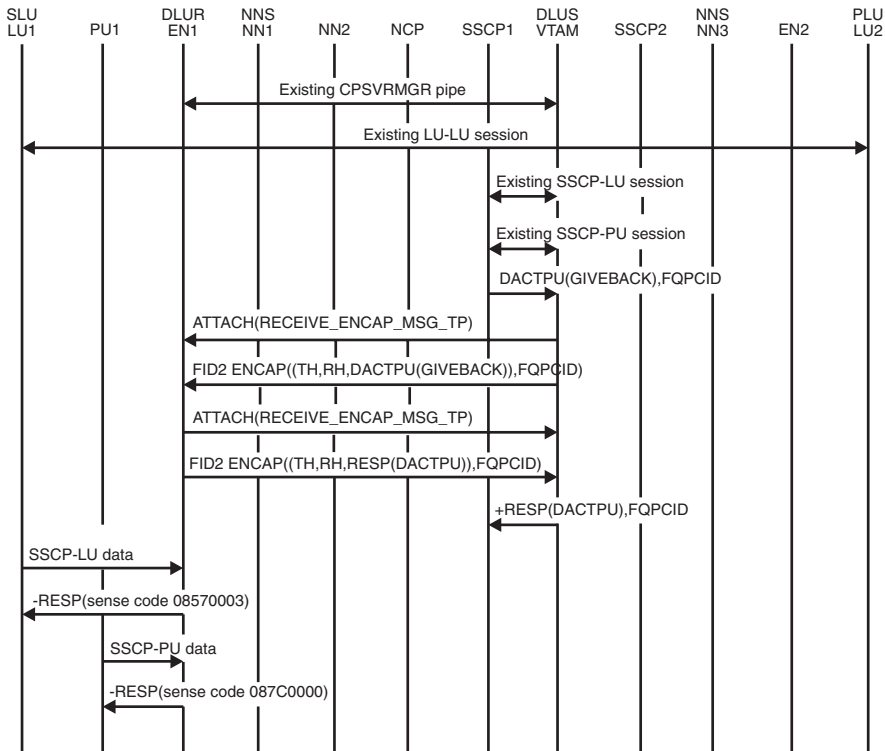


Figure 251. Giveback SSCP-PU/SSCP-LU session deactivation (ANS=CONT)

APPN PLU-initiated LU-LU session to a dependent SLU

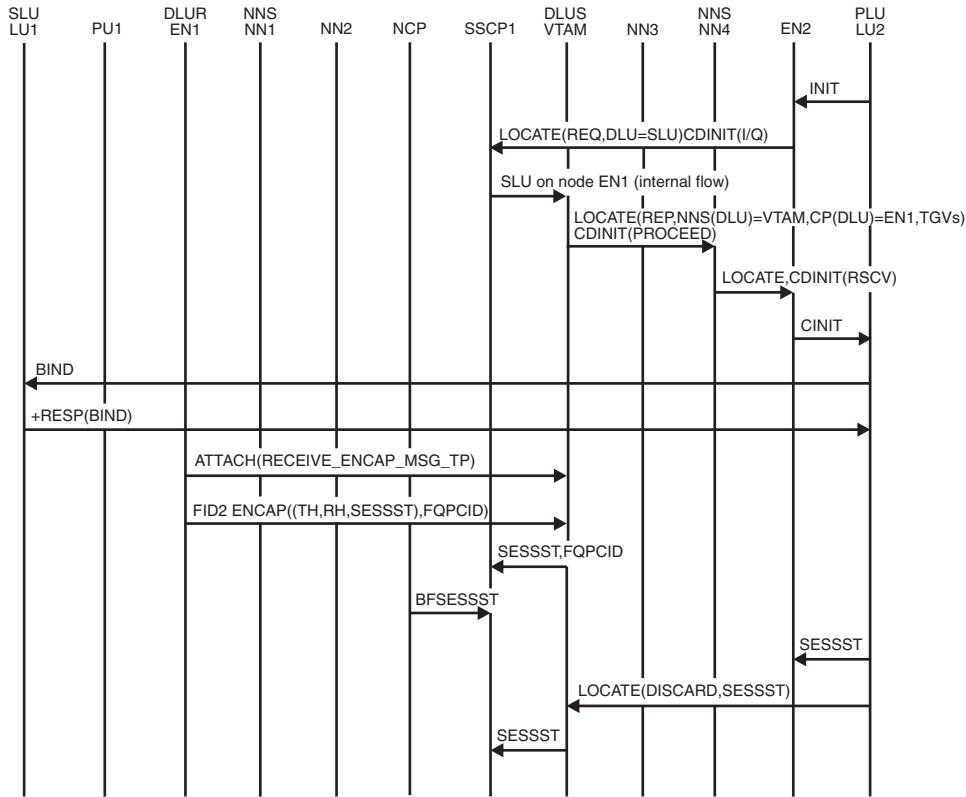


Figure 252. APPN PLU-initiated LU-LU session to a dependent SLU

Note: The transmission group (TG) vectors of the end node-dependent LU requester are provided by previous TG vector registration over the CPSVRMGR pipe.

USS SLU-initiated LU-LU session to APPN PLU

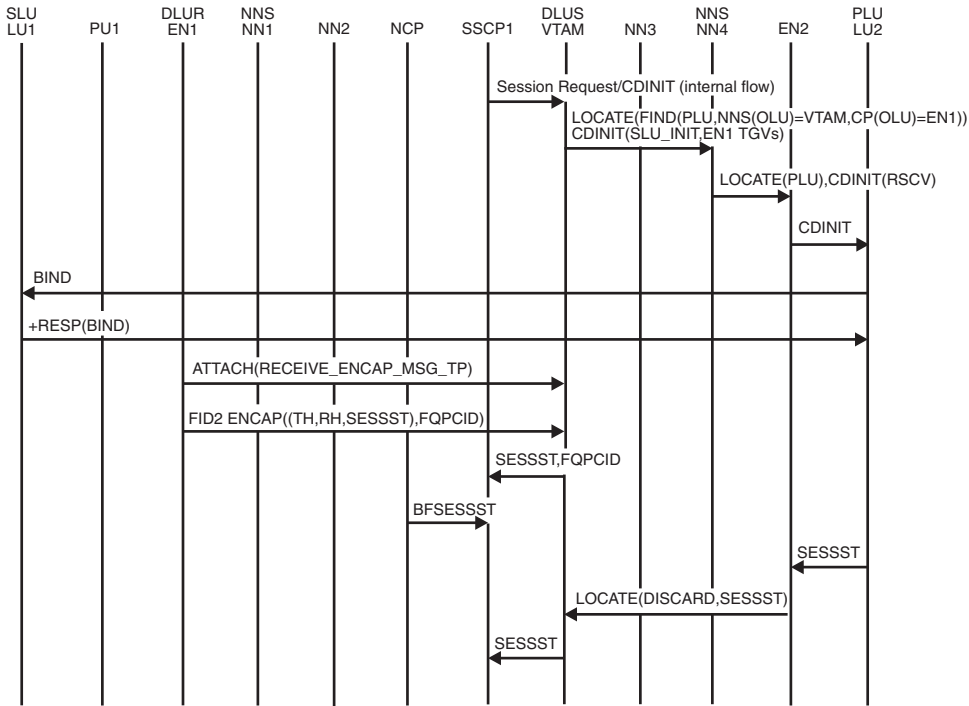


Figure 253. USS SLU-initiated LU-LU session to APPN PLU

Note: The transmission group (TG) vectors of the end node-dependent LU requestor are provided by previous TG vector registration over the CPSVRMGR pipe.

USS SLU-initiated LU-LU session to subarea PLU

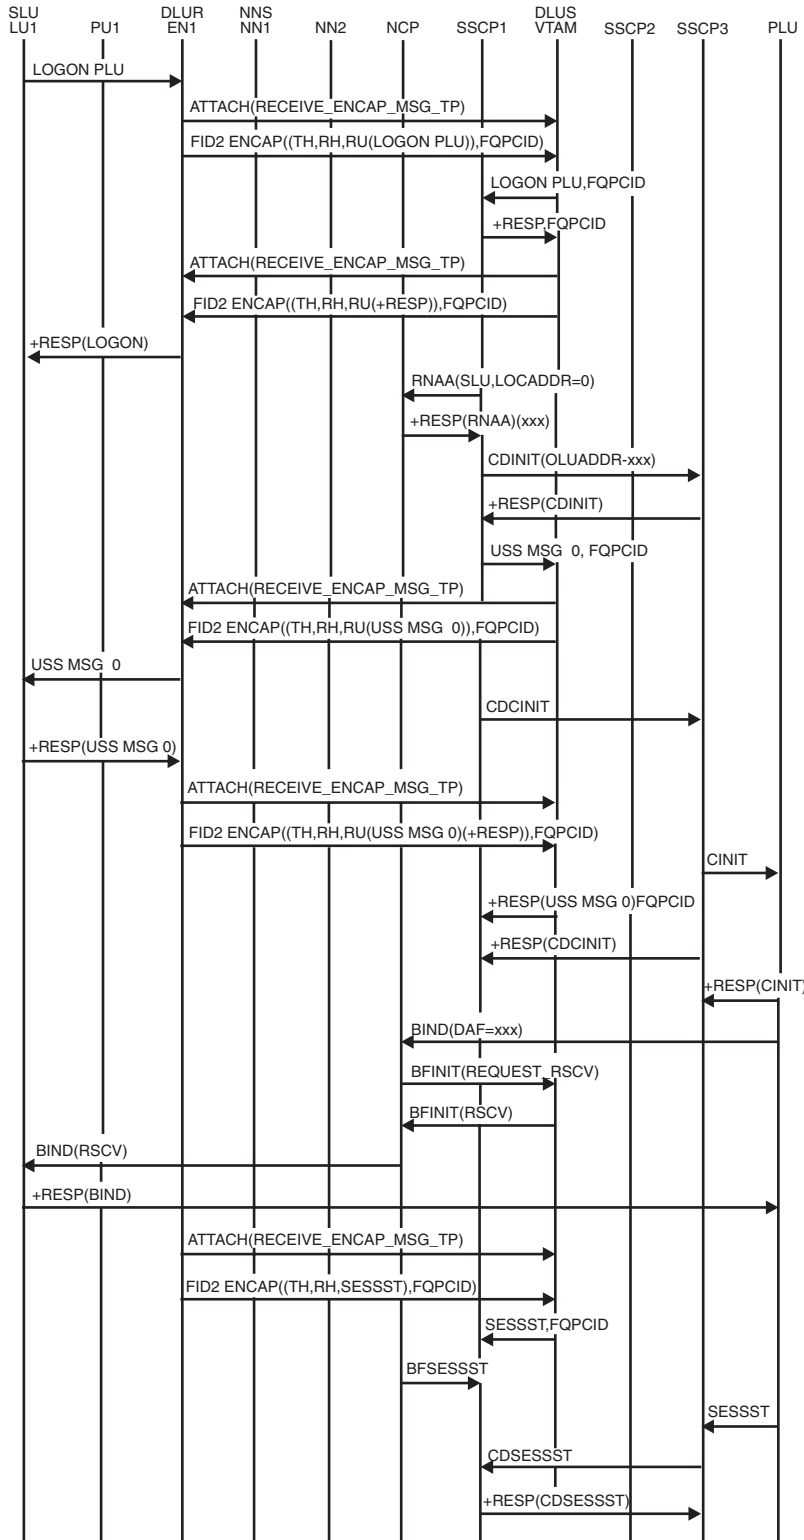


Figure 254. USS SLU-initiated LU-LU session to subarea PLU

USS flows for LU-LU session termination

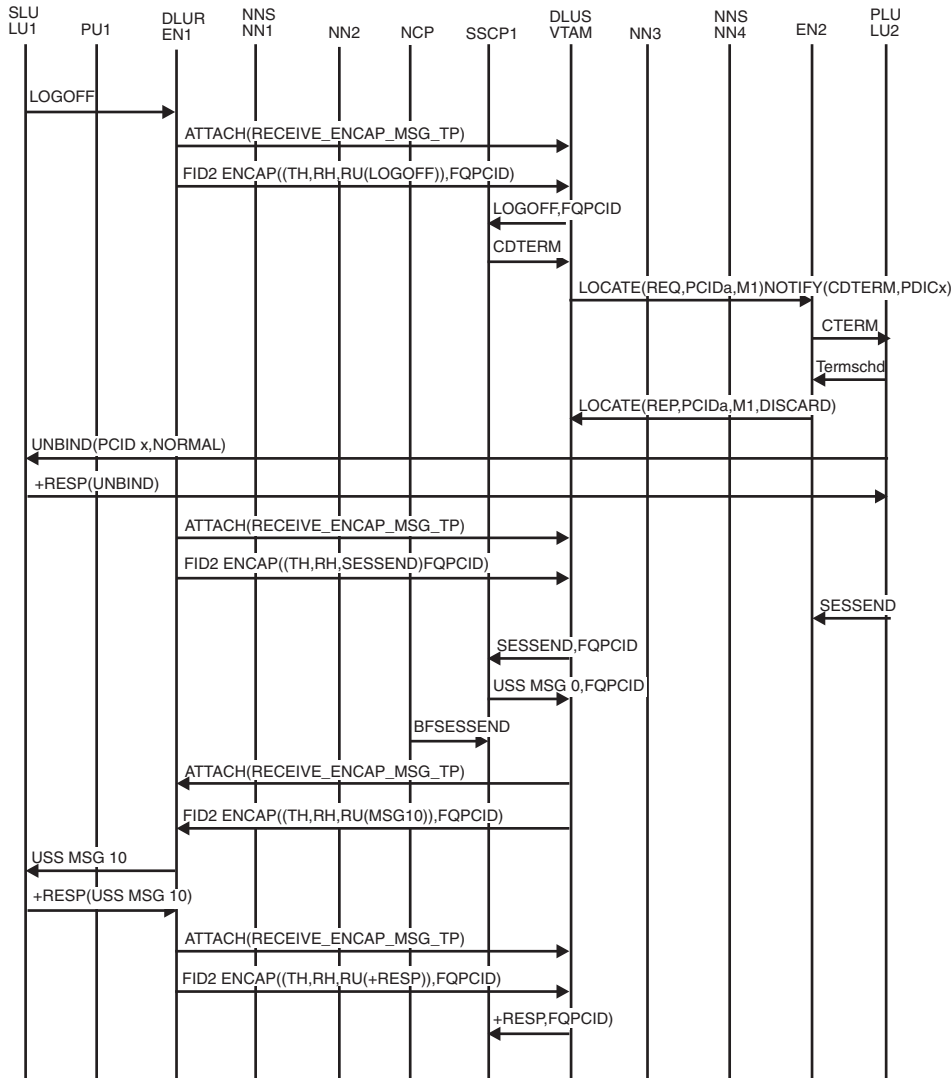


Figure 255. USS flows for LU-LU session termination

Note: The transmission group (TG) vectors of the end node-dependent LU requestor are provided by previous TG vector registration over the CPSVRMGR pipe.

Cross subnetwork flows

Figure 256 on page 606 through Figure 260 on page 611 show the flow of requests and responses between dependent logical unit requestors and servers across subnetworks.

Several abbreviations are used in these flow diagrams.

DSL DLUS-served LU

DSR DLUR search required

ISB Internet search bit
 OCR Owing CP respond

PLU-initiated session with DLUS and DLUR within different subnetworks, PLU Is through the subarea

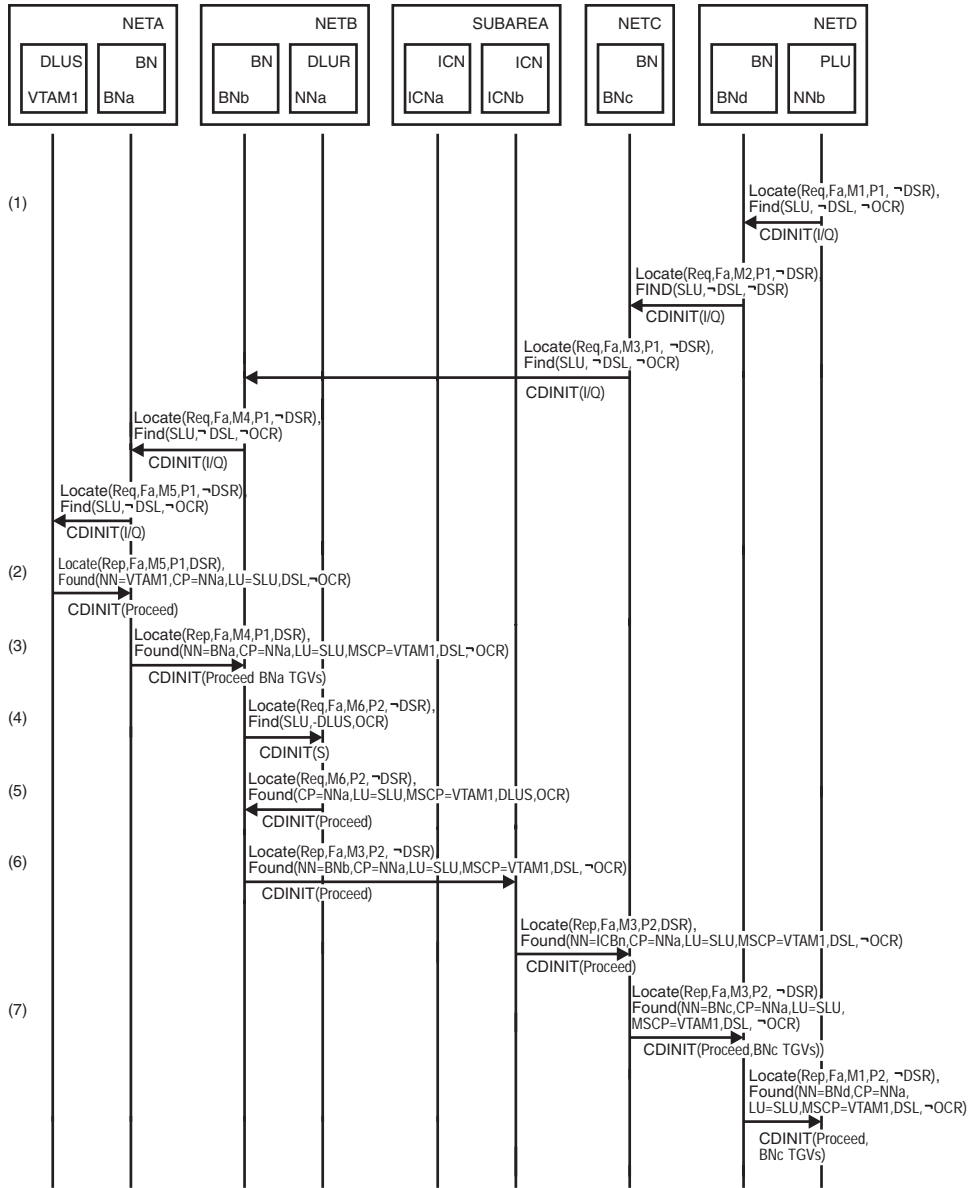


Figure 256. PLU-initiated search with DLUS and DLUR within different subnetworks, PLU through the subarea (part 1 of 2)

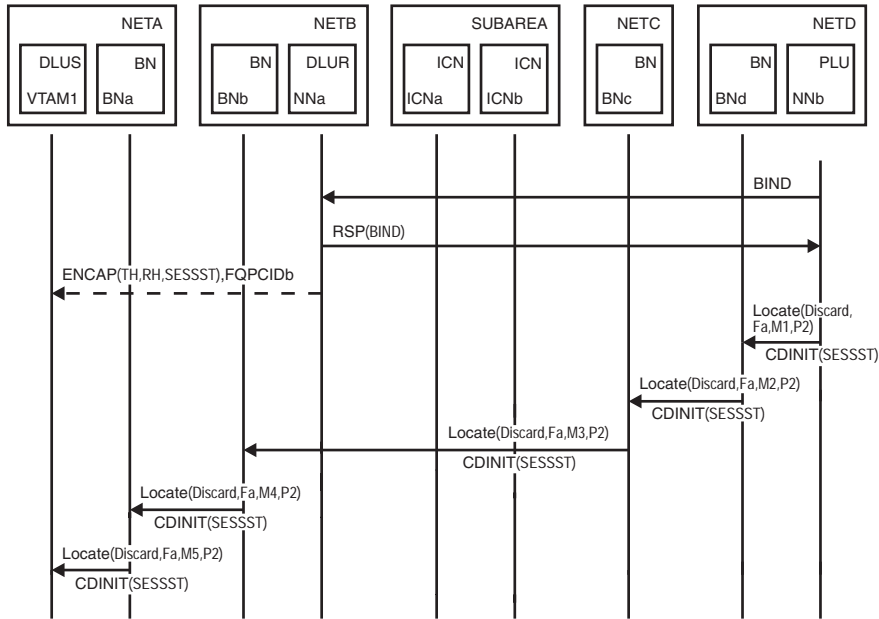


Figure 257. PLU-initiated search with DLUS and DLUR within different subnetworks, PLU through the subarea (part 2 of 2)

1. NNb initiates a search to locate the SLU. NNb has the location of the SLU cached, and the hierarchy indicates that BNd is the NNS(SLU). NNb sends a directed Locate to BNd to verify the location of the SLU and to obtain the SLU TGVs.
 BNd initiates a search to locate the SLU. The search ultimately reaches VTAM1. Neither the DSL nor DSR indicators will be set within the Find GDS variable on the search sent by NNb.
2. The OCR indicator is not set, so VTAM1 responds to the Locate. Because NNa is within a different subnet than VTAM1, and the PLU is also nonnative, VTAM1 will set the DSR indicator on the Locate reply. VTAM1 will also set the DSL indicator on the Locate reply because the SLU is a DLUS-served LU.
3. BNa caches the location of the SLU as being on NNa with VTAM1 as the NNS. Because the OCR indicator is not set on the reply, BNa does not set the OCR indicator within the cache entry. BNa then modifies the Found resource hierarchy to indicate itself as the NNS(DLU) and VTAM1 as the Management Services Control Point (MSCP). BNa also adds its own TGVs to the Locate reply.
4. BNb caches the location of the SLU as being on NNa with BNa as the NNS. Because the OCR indicator is not set on the reply, BNb does not set the OCR indicator within the cache entry.
 Because the DSR and DSL indicators are set on the reply and BNb is returning a reply to a non-Border Node, BNb must obtain the SLU TGs to be included on the Locate reply that will be returned to the NNS(PLU).
 To obtain the TGVs, BNb initiates a Locate search to find the SLU. This search will be a PLU-init Search-Only. The OCR indicator is set, requesting that the DLUR node respond to the Locate request. The DSR indicator will not be set, Because BNb is the node that is performing the extra Locate search.

5. The OCR indicator is set, so NNa responds to the Locate. Because the SLU is a DLUS-served LU, NNa sets the DSL indicator on the reply. When building the reply, NNa will include a CV X'40' that includes the DLUS node CP name.
6. BNb caches the location of the SLU as being on NNa. Because both the DSL and OCR indicators are set on the Locate reply, BNb sets the OCR indicator within the cache entry.

BNb modifies the Found resource hierarchy to indicate itself as the NNS(DLU). BNb removes the BNa TGVs from the reply. Because NNa is a network node, there are no TGVs to add to the Locate reply that is forwarded to ICNa.

7. Because the DSR indicator is not set on the reply, neither BNc nor BNd will submit an extra Locate search to obtain the SLU TGVs even though each is closer to the PLU than BNb.

When NNb receives the Locate reply, NNb calculates an RSCV and sends the BIND to NNa. This establishes the session between the PLU and the SLU.

PLU-initiated session with DLUS and PLU in one subnetwork and DLUR in another

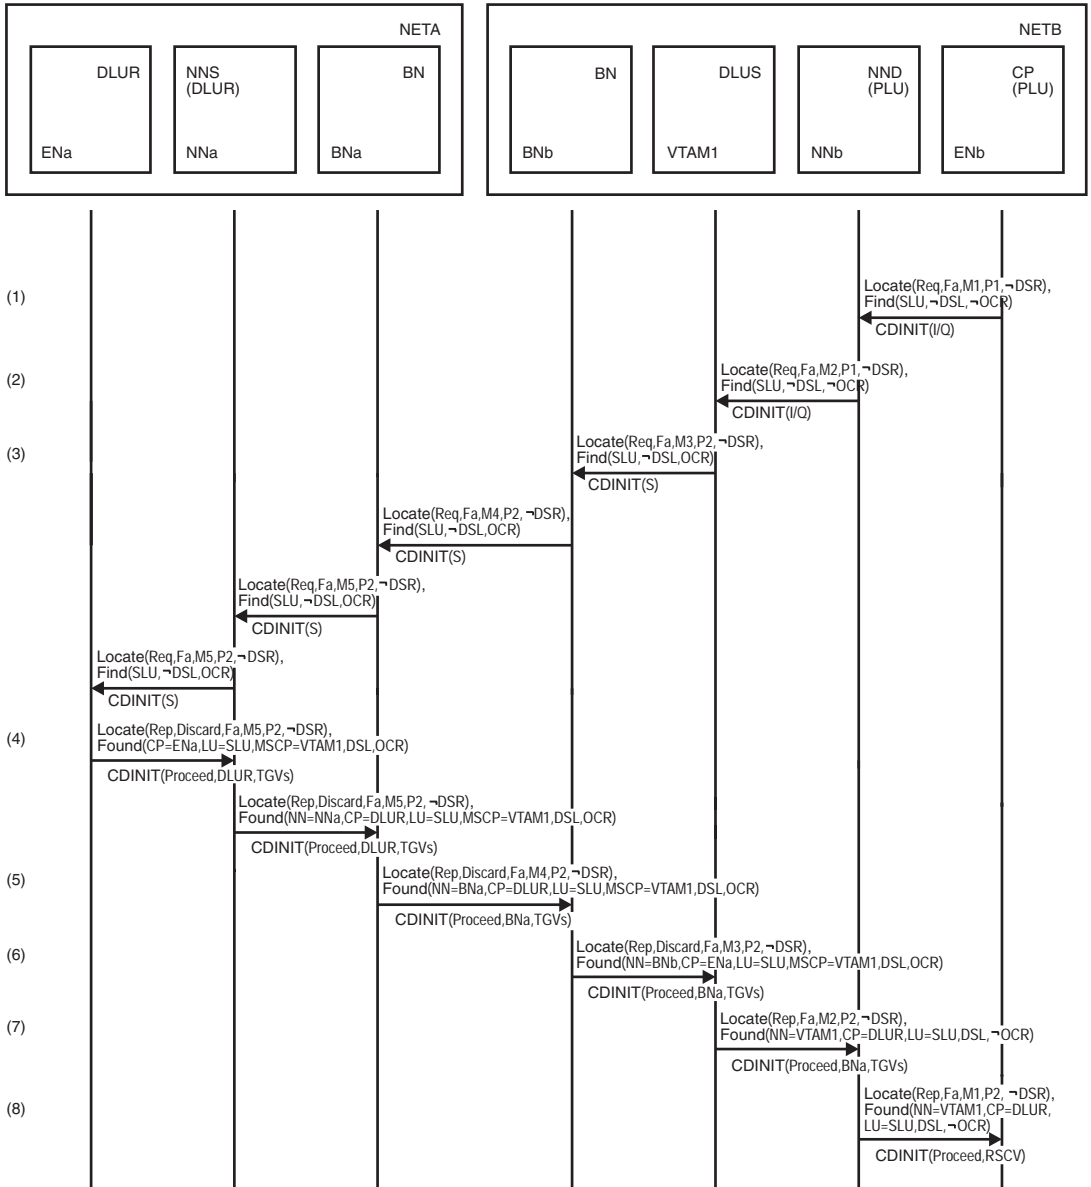


Figure 258. PLU-initiated session with DLUS and PLU in same subnetwork and DLUR in another (part 1 of 2)




Figure 259. PLU-initiated session with DLUS and PLU in same subnetwork and DLUR in another (part 2 of 2)

1. ENb initiates a search to locate the SLU. The search is forwarded to the ENb network node server. Neither the DSL, DSR, nor OCR indicators will be set within the Find GDS variable.
2. NNb initiates a search to locate the SLU. NNb has the location of the SLU cached, and the hierarchy indicates that VTAM1 is the NNS(SLU). NNb sends a directed Locate to VTAM1 to verify the location of the SLU and to obtain the SLU TGVs.
3. The OCR indicator is not set, so VTAM1 responds to the locate. Because ENa is in a different subnet than VTAM1, no endpoint TGVs were reported over the CPSVRMGR pipe. Because NNb is within the same APPN subnet as VTAM1, VTAM1 must obtain the endpoint TGVs to be included in the Locate reply returned to NNb. To obtain the endpoint TGVs, VTAM1 initiates a new Locate search to find the SLU. This search will be a PLU-init Search-Only. The OCR indicator will be set, requesting that the DLUR node respond to the Locate request. The DSR indicator will not be set, because VTAM1 is the node that is performing the extra Locate search. When initiating the Locate search, a new PCID modifier slot will be allocated and the PRN will be incremented. This will allow the Locate search to appear as a new search within both the APPN and subarea networks. When searching their caches for the SLU, VTAM1, BNb, and BNa all look for entries where the OCR indicator is set. These entries will allow the nodes to route the Locate search to the DLUR node instead of the DLUS node.
4. The OCR indicator is set, so ENa responds to the Locate. Because the SLU is a DLUS-served LU, ENa sets the DSL indicator. When building the reply, ENa will include a CV X'40' which includes the DLUS node CP name.
5. BNa caches the location of the SLU as being on the DLUR with NNa as the network node server. Because both the DSL and OCR indicators are set on the search reply, BNa sets an OCR indicator within the cache entry. BNa then modifies the Found resource hierarchy to indicate itself as the NNS(DLU). BNa also replaces the DLUR TGVs with its own TGVs before forwarding the Found to BNb.
6. BNb caches the location of the SLU, with BNa as the NNS. As with BNa, an OCR indicator is saved with the cache entry. BNb then modifies the resource hierarchy so that it appears as the NNS(DLU). The Locate reply is then forwarded to VTAM1.
7. VTAM1 caches the location of the SLU. As with BNa and BNb, VTAM1 sets the OCR indicator within the cache entry. VTAM1 then replies to the Locate search request received from NNb. Because the SLU is a DLUS-served LU, VTAM1 alters the Found hierarchy to indicate that VTAM1 is the NNS(DLU) and ENa is the CP(DLU). VTAM1 then removes the TGVs returned on the Locate reply received from BNb and places the TGVs on the Locate reply that it is constructing. Because the SLU is a DLUS-served resource, the DSL indicator is set. However, the DSR indicator is not set because VTAM1 has already obtained the correct DLUR TGVs.
8. NNb calculates an RSCV using the endpoint TGVs returned by VTAM1. NNb returns the RSCV to ENb on the Locate reply. ENb then BINDs the session between the PLU and the SLU.

SLU-initiated session with DLUS and DLUR within different subnetworks

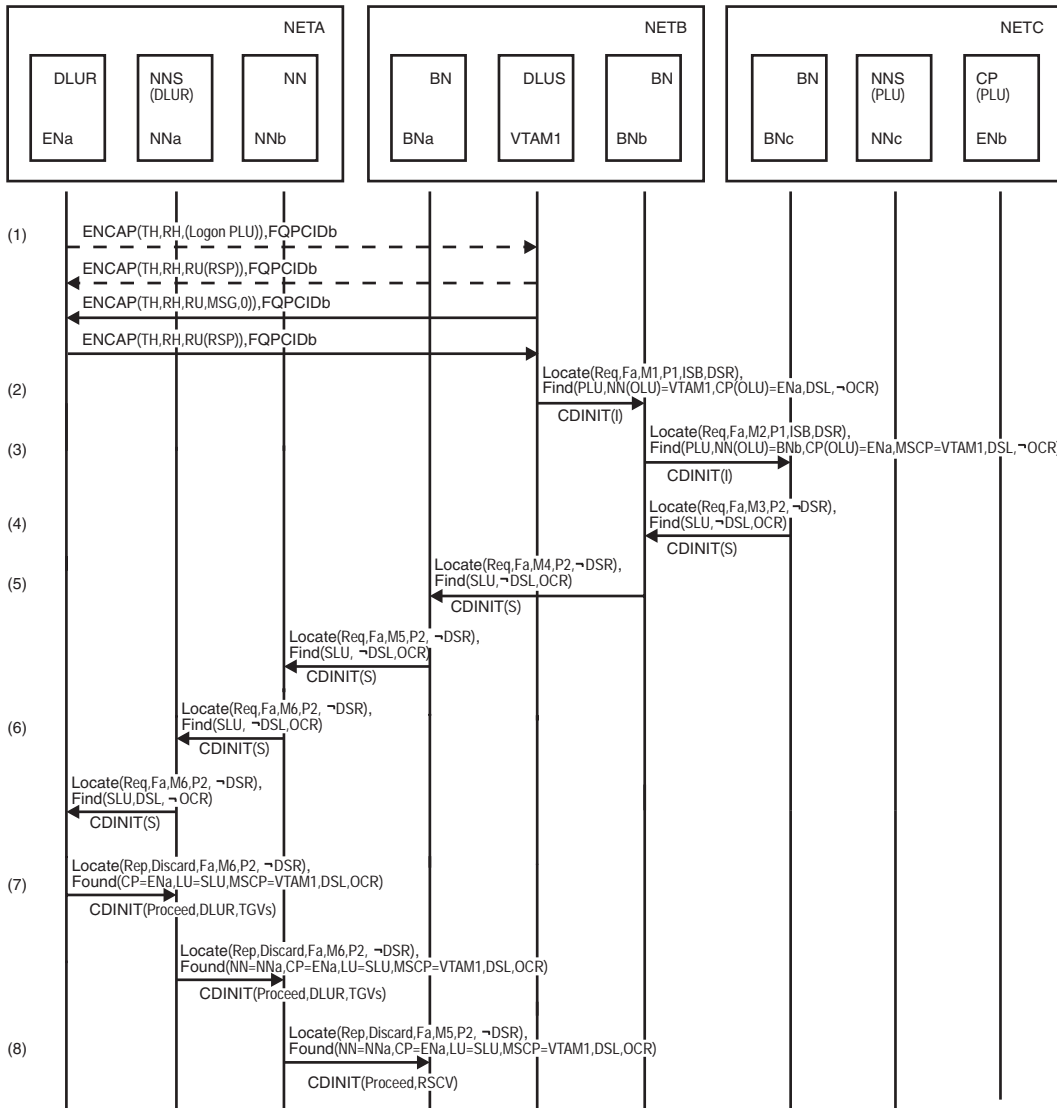


Figure 260. SLU-initiated session with DLUS and DLUR within different subnetworks (part 1 of 2)

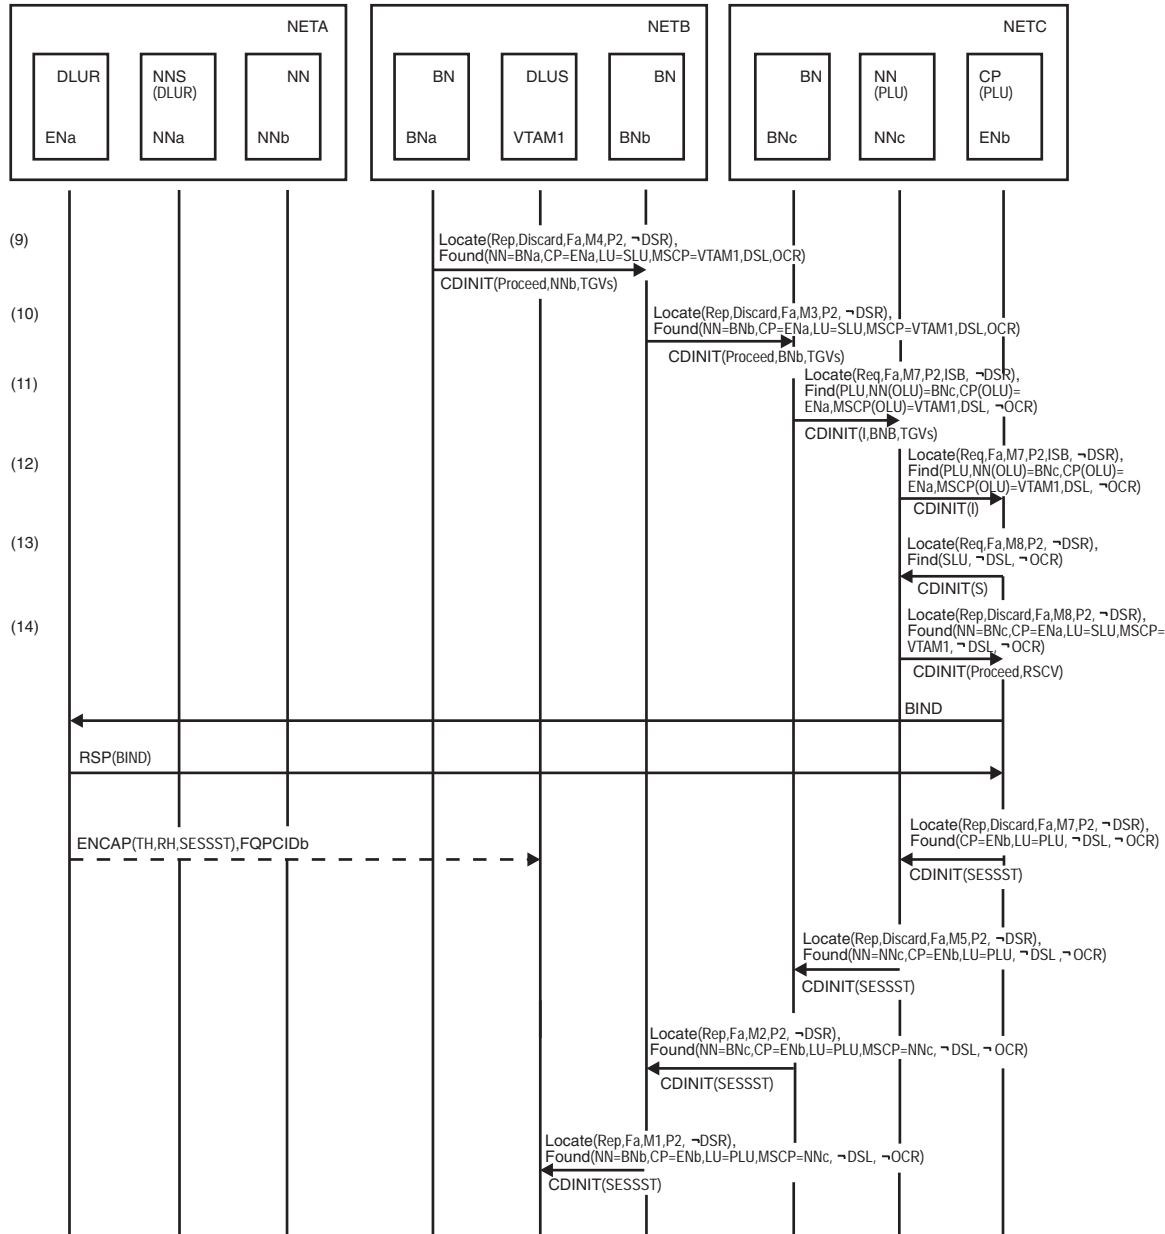


Figure 261. SLU-initiated session with DLUS and DLUR within different subnetworks (part 2 of 2)

1. The SLU initiates a Logon, which ENa encapsulates on the CPSVRMGR pipe and sends to VTAM1. VTAM1 sends a response to the Logon request followed by USS MSG0. ENa sends a response to USS MSG0 to VTAM1.
2. before initiating a search to locate the PLU, VTAM1 examines its cache. The PLU is located and is cached as being accessible through BNb. Because ENa is in a different subnet than VTAM1, no endpoint TGVs were reported over the CPSVRMGR pipe. Because the PLU is also within a different APPN subnet, VTAM1 will set the DSR indicator on the Locate request. VTAM1 will also set the DSL indicator on the Locate request because the SLU is a DLUS-served LU.

3. BNb caches the location of the SLU as being on ENa with VTAM1 as the NNS. The OCR indicator will not be set in the cache entry. BNb has the PLU cached as being accessible through BNc.
BNb modifies the Find resource hierarchy to indicate itself as the NNS(OLU) and adds a CV X'40' with VTAM1 as the MSCP. BNb also adds its endpoint TGs to the Locate search request and forwards the request to BNc.
4. BNc caches the location of the SLU as being on BNb. The OCR indicator is not set within the cache entry.
BNc has the location of the PLU cached as being within the native subnet. Because both the DSL indicator and the DSR indicator are set on the Locate request, BNc must obtain the SLU TGVs to be included on the Locate search request.
To obtain the endpoint TGVs, BNc initiates a Locate search to find the SLU. This search will be a PLU-init Search-Only. The OCR indicator is set, requesting that the DLUR node respond to the Locate request. The DSR indicator will not be set, because BNc is the node which is performing the extra Locate search.
5. BNb receives the new Locate search from BNc. BNb finds a cache entry for the SLU with the OCR indicator set. The entry indicates that BNa is the NNS(SLU), so BNb forwards the Locate search to BNa.
BNa also finds a cache entry for the SLU with the OCR indicator set. The cache entry for BNa indicates that the search should be forwarded to NNb.
6. NNb receives the Locate search from BNa. As part of its search logic, NNb will send either a directed Locate search to NNa or will perform a network broadcast that will ultimately reach NNa. Either way, the Locate search will be forwarded to NNa and, ultimately, ENa.
7. The OCR indicator is set, so ENa responds to the Locate. Because the SLU is a DLUS-served LU, ENa sets the DSL indicator on the reply. When building the reply, ENa will include a CV X'40' that includes the DLUS node CP name.
8. NNb caches the location of the SLU as being on ENa, with NNa as the NNS. NNb then calculates an RSCV for the Bind route between BNa and ENa and returns the RSCV to BNa.
9. BNa caches the location of the SLU as being on ENa, with NNb as the NNS. Because both the DSL and OCR indicators are set on the Locate reply, BNa sets the OCR indicator within the cache entry.
BNa modifies the Found resource hierarchy to indicate itself as the NNS(DLU). BNa then removes the RSCV from the Locate reply and places the NNb TGVs on the reply. The reply is then returned to BNb.
10. BNb caches the location of the SLU as being on ENa with BNa as the NNS. Both the DSL and OCR indicators are on the Locate reply, and BNb sets the OCR indicator on the reply. BNb then modifies the resource hierarchy in the reply to indicate itself as the NNS(DLU). BNb also replaces the NNb TGs with its own TGs and then forwards the reply to BNc.
11. BNc caches the location of the SLU as being on ENa with BNb as the NNS. Because both the DSL and OCR indicators are set on the Locate reply, BNc sets the OCR indicator within the cache entry.
BNc then takes the endpoint TGVs that were returned on the Locate reply that was just received from BNb and places the TGVs on the Locate request that was received from BNb. BNc then modifies the resource hierarchy on the request to indicate itself as the NNS(OLU). The search request is then sent to NNc.
12. NNc forwards the search request to ENb.

13. Because this is a SLU-init search request and an RSCV was not present on the Locate request, ENb initiates a PLU-init search with the SLU as the target.
14. NNC correlates the PLU-init search request with the outstanding SLU-init search request. Using the information that was provided on the original SLU-init search, NNC calculates an RSCV and returns it to ENb on the Locate reply. ENb then BINDs the session between the PLU and the SLU.
Because NNC creates the Locate reply instead of VTAM1, neither the DSL indicator nor the DSR indicator will be set on the reply. This can be contrasted to the original SLU-init request in which the DSL indicator was set.

High-Performance Routing flows

Figure 262 on page 615 through Figure 267 on page 620 show network flows for the high-performance routing function. For more information, see *z/OS Communications Server: SNA Network Implementation Guide*.

Index of High-Performance Routing flows

Table 44 lists the flows illustrated here.

Table 44. Index of High-Performance Routing flows

Flow	Page
Rapid-transport protocol (RTP) connection over portion of session path	Figure 264 on page 617
Rapid-transport protocol (RTP) across composite nodes with a T2.1 connection through NCP	Figure 265 on page 618
Rapid-transport protocol (RTP) across composite nodes with a T2.1 connection through VTAM	Figure 266 on page 619
Rapid-transport protocol (RTP) across composite nodes with a virtual-route-based transmission group, NCP does ANR routing	Figure 267 on page 620
Rapid-transport protocol (RTP) across composite nodes with a virtual-route-based transmission group, VTAM does ANR routing	Figure 268 on page 621
Two rapid-transport protocol (RTP) nodes with a T2.1 connection	Figure 262 on page 615
Two rapid-transport protocol (RTP) nodes with a virtual-route-based transmission group	Figure 263 on page 616

Two Rapid-Transport Protocol (RTP) nodes with a T2.1 connection

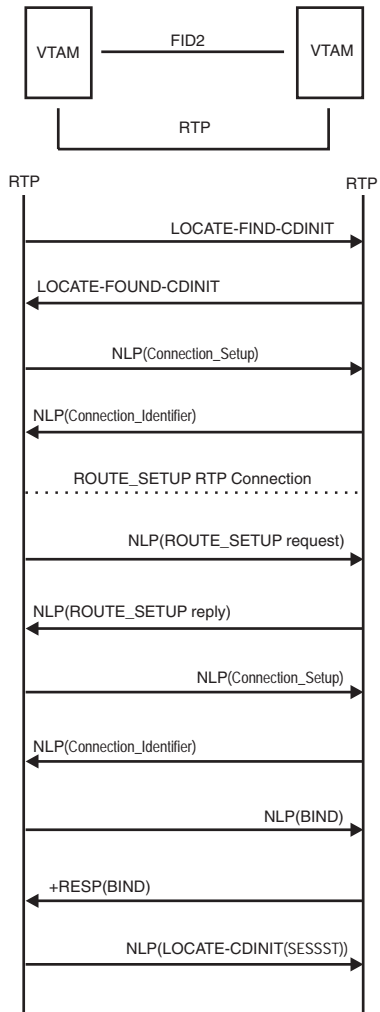


Figure 262. An example of two Rapid-Transport Protocol (RTP) nodes with a T2.1 connection

Two Rapid-Transport Protocol (RTP) nodes with a virtual-route-based transmission group

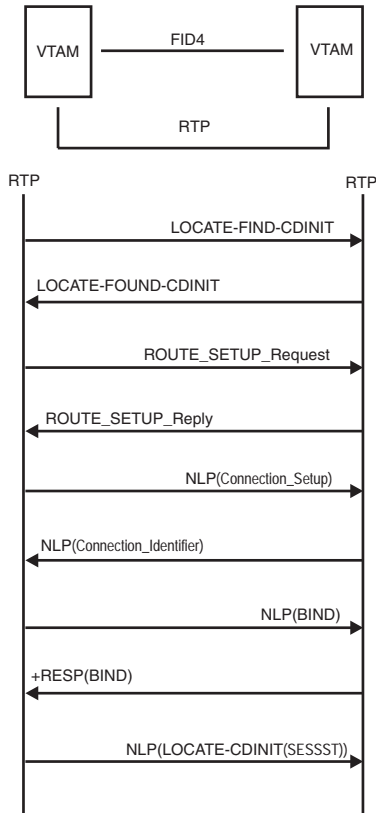


Figure 263. Two Rapid-Transport Protocol (RTP) nodes with virtual-route-based transmission group

Rapid-Transport Protocol (RTP) connection over portion of session path

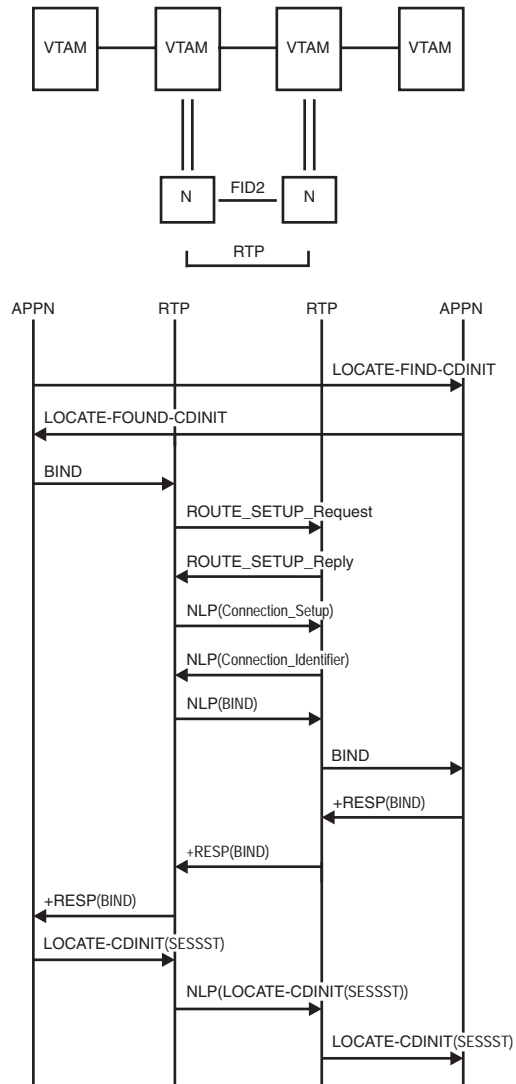


Figure 264. Rapid-Transport Protocol (RTP) connection over portion of session path

Rapid-Transport Protocol (RTP) across composite nodes with T2.1 connection through NCP

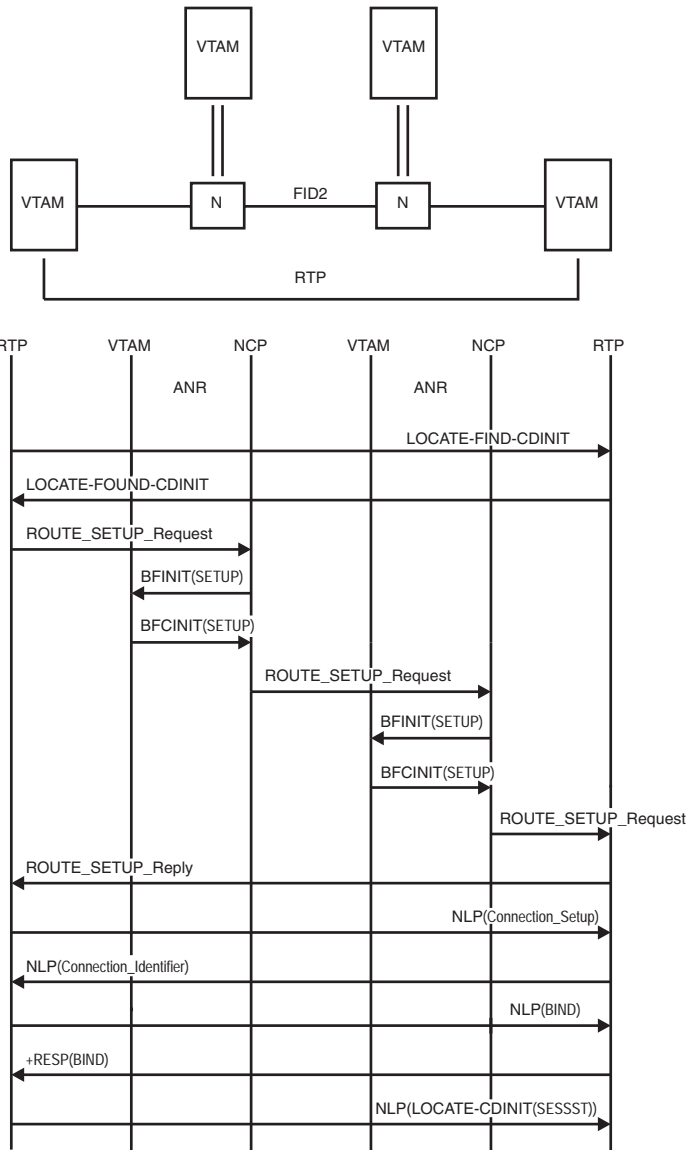


Figure 265. Rapid-Transport Protocol (RTP) across composite nodes with T2.1 connection through NCP

Rapid-Transport Protocol (RTP) across composite nodes with T2.1 connection through VTAM

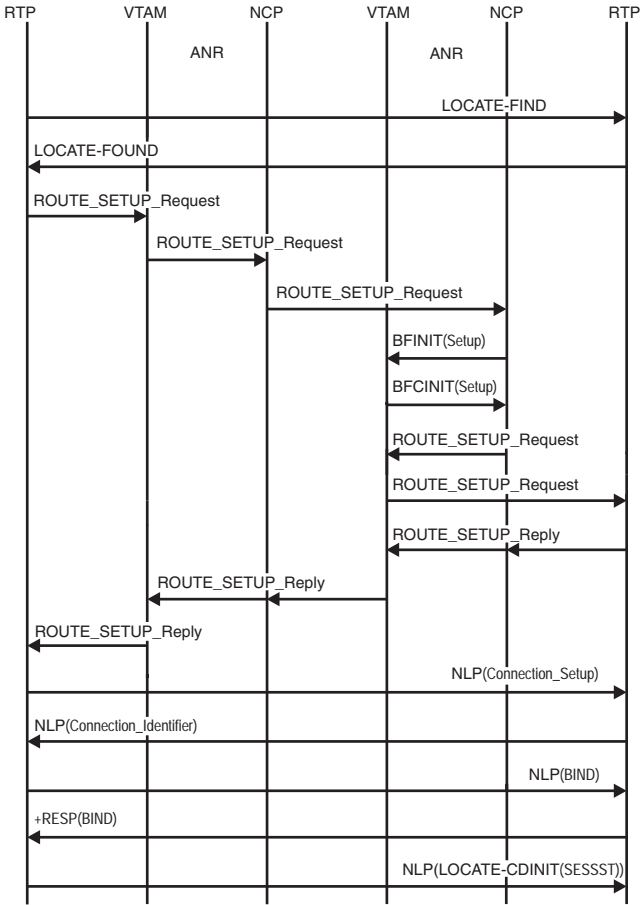
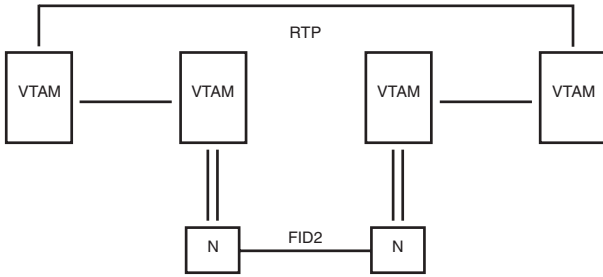


Figure 266. Rapid-Transport Protocol (RTP) across composite nodes with T2.1 connection through VTAM

Rapid-Transport Protocol (RTP) across composite nodes with a virtual-route-based transmission group, NCP does ANR routing

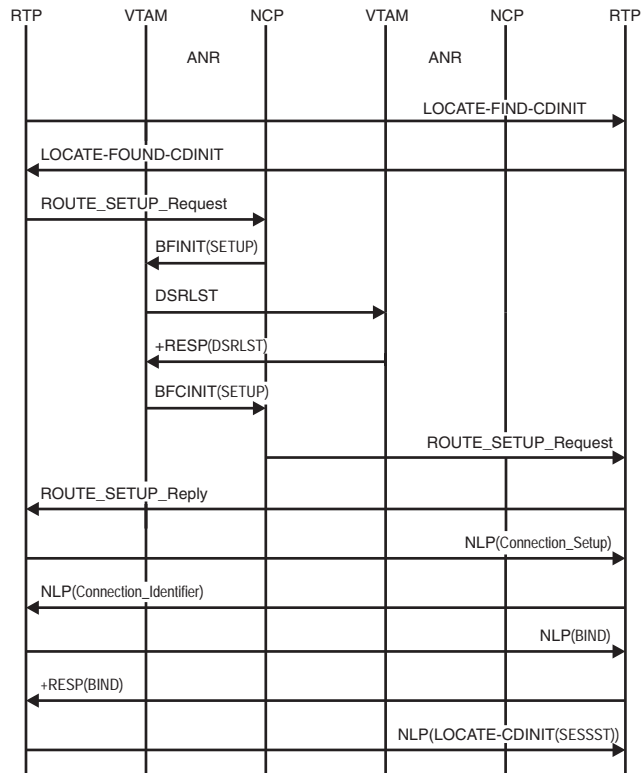
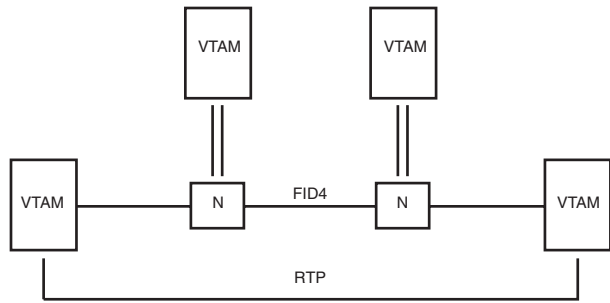


Figure 267. Rapid-Transport Protocol (RTP) across composite nodes with a virtual-route-based transmission group, NCP does ANR routing

Rapid-Transport Protocol (RTP) across composite nodes with a virtual-route-based transmission group, VTAM does ANR routing

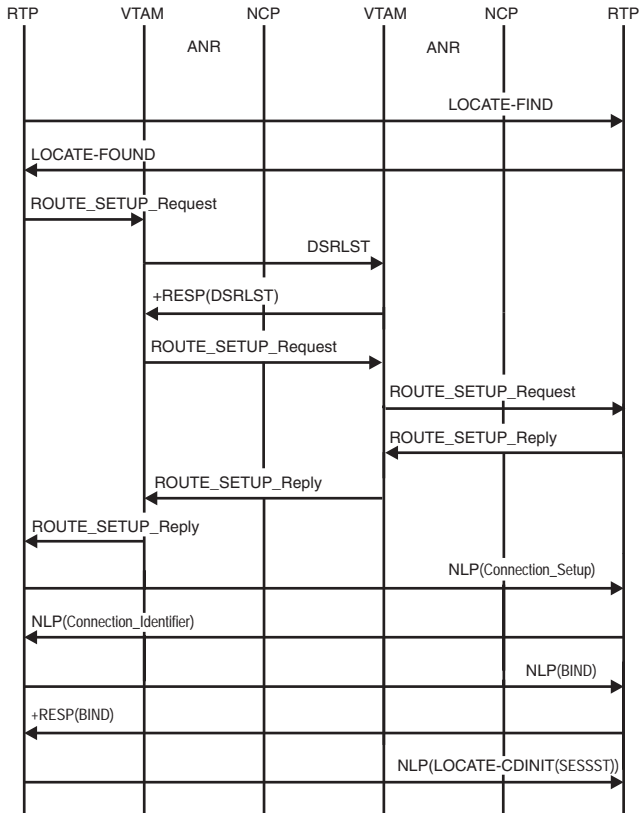
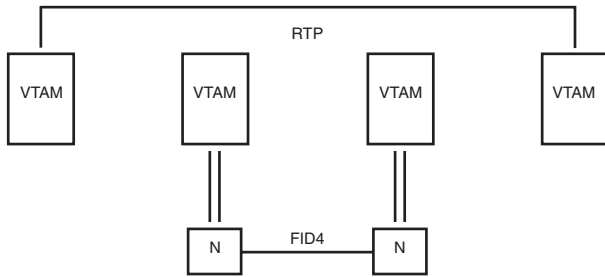


Figure 268. Rapid-Transport Protocol (RTP) across composite nodes with a virtual-route-based transmission group, VTAM does ANR routing

Appendix D. Control point/control block (CPCB) operation codes

Many processes of the VTAM program are represented by RUPE, NCSPL, DLRPL, CPCB, PPL, NOSPL, MFT, FILTR, or TQE work elements. Each of these work elements contains a prefix called a CPCB at the beginning of the control block. The CPCB prefix contains a field called the CPCB operation code (CPCBOPC), which provides an indication of the type of VTAM process represented by the work element.

The CPCBOPCs are mapped by the ISTCPKCB regarding the data area. The CPCBOPC field is 4 bytes long and contains a category byte (CPCBCAT), followed by a 3-byte specific operation code field (CPCBFMH).

The contents of the CPCB operation code category are summarized in the following table:

Category (hex)	Meaning
00	Operator command
01	Dump/load/restart request
02	I/O purge request
03	Timer management request
04	Unformatted RU
08	Function management data RU
09	Network control RU
0A	Data flow control RU
0B	Session control RU
0C	Function management data access method RU
0D	Network control access method RU
0E	Data flow control access method RU
0F	Session control access method RU
10	VTAM topology agent services request
FC	Function management data access method interprocess signal
FD	Network interprocess signal

Table 45 summarizes the possible values of the CPCB operation code field.

Note:

1. For each CPCB operation code listed, the character string representation that may appear in VTAM operator messages is given, along with the function of the associated work element.
2. Internal codes that are used only by the product support organization to assist in internal flow diagnosis are not included in this list.

Table 45. Control point/control block operation codes (CPCBOPC)

CPCBOPC	Message display	Function
00010000	VARY	VARY Command
00010001	VARY ACT	VARY Activate

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
00010002	VARY INACT	VARY Deactivate
00010004	VARY LOGON	VARY LOGON
00010005	VARY ACT	VARY Activate, LOGON Parameter Specified
00010008	VARY DRDS	VARY DRDS
00010012	VARY INACT	VARY Deactivate Immediate
00010032	VARY INACT	VARY Deactivate Immediate (Internal)
00010040	VARY ANS	VARY ANS
00010041	VARY ACT	VARY Activate, ANS Parameter Specified
00010045	VARY ACT	VARY Activate, ANS and LOGON Parameters Specified
00010080	VARY HGUP	VARY HANGUP
00010100	VARY PATH	VARY PATH=USE
00010200	VARY PATH	VARY PATH=NOUSE
00010400	VARY INOP	VARY INOP
00010802	FORCE DEAC	Force Deactivate
00011002	FORCE REAC	Force Reactivate
00011802	INACT GVBK	VARY Deactivate Giveback
00012000	VARY ACQ	VARY Acquire
00012001	VARY ACQ	VARY Activate, ACQ Parameter Specified
00012002	INACT SON	Deactivate (Session Outage Notification)
00012005	VARY ACQ	VARY Activate, ACQ and LOGON Parameters Specified
00014000	VARY REL	VARY Release
00014010	REL IMMED	VARY Release Immediate
00014012	REL GVBK	VARY Release Giveback
00018000	VARY DIAL	VARY DIAL
00018004	VARY NOLOG	VARY NOLOGON
00020000	F EXIT	MODIFY EXIT CMD
00020001	F DUMP	MODIFY DUMP
00020002	F ENCR	MODIFY ENCR
00020003	F TGP	MODIFY TGP
00020004	F CHANGE	MODIFY CHANGE
00020005	F SECURITY	MODIFY SECURITY
00020010	F TABLE	MODIFY TABLE
000200A0	F ACT NCTR	MODIFY TRACE, TYPE=NETCTLR

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
000200C0	F DACTNCTR	MODIFY NOTRACE, TYPE=NETCTLR
00020120	F ACT GPT	MODIFY Activate Generalized PIU Trace
00020140	F DACT GPT	MODIFY Deactivate Generalized PIU Trace
00020180	F LOAD ADD	MODIFY LOAD ADD
00020200	F CDRM	MODIFY CDRM
00020240	SETTIME	SETTIME Cancel
00020280	F LOAD REP	MODIFY LOAD REPLACE
00020401	F DUMP TRN	Transfer DUMP (NCP)
00020408	F DR MOVE	MODIFY DR MOVE
00020480	F LOAD PRG	MODIFY LOAD PURGE
00020800	F LINEDEF	MODIFY LINEDEF
00020801	F DUMP PGN	PURGE DUMP (NCP)
00020808	F DR DEL	MODIFY DR DELETE
00020820	F ALTRACE	MODIFY Activate Line Trace
00020840	F DLTRACE	MODIFY Deactivate Line Trace
00020880	F LOAD CAN	MODIFY LOAD CANCEL
00021001	F DUMP CSP	MODIFY DUMP (CSP)
00021080	F LOAD SET	MODIFY LOAD ACTION=SETTIME
00021801	F DUMP PGC	PURGE DUMP (CSP)
00022001	F DUMP MOS	MODIFY DUMP (MOSS)
00022080	F LOAD REN	MODIFY LOAD ACTION=RENAME
00022801	F DUMP PGM	PURGE DUMP (MOSS)
00024001	F DUMP DYN	MODIFY Dump (Dynamic)
00024020	F ACT SIT	MODIFY Activate SIT Trace
00024040	F DACT SIT	MODIFY Deactivate SIT Trace
00025000	F RTP	MODIFY RTP
00028001	F DUMP DYN	MODIFY Dump (Dynamic-CH)
00028820	F ACT TG	MODIFY Activate TG Trace
00028840	F DACT TG	MODIFY Deactivate TG Trace
0002C000	F ALSLIST	MODIFY Adjacent Link Station List
0002D001	F RESOURCE	MODIFY RESOURCE
0002D000	F DEFAULTS	MODIFY DEFAULTS
0002E001	F DIR DEL	MODIFY DIR DELETE
0002E002	F DIR UDP	MODIFY DIR UPDATE
0002E003	F QUERY	MODIFY QUERY
00030001	SOFT INOP	SOFT INOP
00030002	SSCP TKOVR	SSCP TAKEOVER
00030004	HARD INOP	HARD INOP
00040000	DISPLAY	Display Internal Commands
00040001	DISPLAY DLRS	DISPLAY DLURs

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
00060001	INT SYNCH	Internal Synchronization Function
00060002	IPL INIT	IPL Initial Request
00060004	IPL TEXT	IPL Text Request
00060008	IPL FINAL	IPL Final Request
00060010	DUMP INIT	Dump Initial Request
00060020	DUMP TEXT	Dump Text Request
00060040	DUMP FINAL	Dump Final Request
00060100	HALT CDLNK	Process Cross-Domain Links During HALT
00061001	REQDMP CSP	Request CSP Dump
00062000	RESET LU	Reset LU
00062001	REQDMP MOS	Request MOSS Dump
00063001	F DUMP TRH	Internal Transfer of Dump Header
00063002	F DUMP TRM	Internal Transfer of Dump Main Storage
00064000	REQLOAD	Request Load
00064001	REQDUMP DY	Request Dynamic Dump Data
00068000	REQDUMP	Request Dump
00080001	DIAL START	Dial Start Request
012B0000	CKPTN	Checkpoint Node Status Function
014B0000	CHKPT	Checkpoint Resource Status Function
01BD0000	CPMSG	Internal WTOR Function
01DD0000	DLR PURGE	Dump/Load/Restart Purge
01EA0000	CPCRYPT	Cryptography Management Function
01EB0000	GETHOSTBNM	Gethostbyname
01ED0000	SELECT VR	Virtual Route Select
02510000	CDRM CLEAR	Clear CDRM-CDRM Session
02520000	CDRM ERP	CDRM ERP Internal Clear
03000000	TIMER REQ	Set Timer Request
04000000	CHAR CODED	Unformatted Request Unit
08010001	CHG TLIMIT	Change Transmission Limit
08010002	CHG NRSPOL	Change Negative Response to Poll Limit
08010003	CHG SESSLM	Change Session Limit
08010004	CHG POLLIM	Change Poll Limit
08010201	CONTACT	Contact
08010202	DISCONTACT	Discontact
08010203	IPL INIT	NC IPL Initial
08010204	IPL TEXT	NC IPL Text
08010205	IPL FINAL	NC IPL Final
08010206	DUMP INIT	Dump Initial
08010207	DUMP TEXT	Dump Text
08010208	DUMP FINAL	Dump Final

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
08010209	RMPO	Remote Power [®] Off
0801020A	ACTLINK	Activate Link
0801020B	DACTLINK	Deactivate Link
0801020E	CONNOUT	Connect Out
0801020F	ABCONN	Abandon Connection
08010211	SCV	Set Control Vector
08010213	NTNMON RPLY	NTUNEMON Reply
08010214	ENT SLOWDN	Enter Slowdown
08010215	EXT SLOWDN	Exit Slowdown
08010216	ACTCONNIN	Activate Connect In
08010217	DACTCONNIN	Deactivate Connect In
08010218	ABCONNOUT	Abandon Connect Out
08010219	ANA	Assign Network Address
0801021A	FNA	Free Network Address
0801021B	REQDISCONT	Request Discontact
08010280	CONTACTED	Contacted
08010281	INOP	Inoperative
08010284	REQCONT	Request Contact
08010285	NSLSA	Network Services Lost Subarea
08010302	ACTTRACE	Activate Trace
08010303	DACTTRACE	Deactivate Trace
08010331	DISP STOR	Display Storage
08010334	RECSTOR	Record Storage
08010381	RECMS	Record Maintenance Statistics
08010382	REC TEST	Record Test Data
08010383	REC TRACE	Record Line Trace Data
08010604	NSPE	Network Services Procedure Error
08010681	INIT SELF	Initiate(Self) Format 0
08010683	TERM SELF	Terminate(Self) Format 0
0812C100	GDS CP_CAP	CP Capabilities
0812C200	GDS TDU	TDU
0812C300	GDS REGSTR	REGISTER Resource
0812C400	GDS LOCATE	LOCATE Resource
0812C500	GDS CDINIT	CDINIT
0812C900	GDS DELETE	DELETE
0812CA00	GDS FIND	FIND
0812CB00	GDS FOUND	FOUND
0812CC00	GDS NOTIFY	NOTIFY
0812CD00	GDS IOCD	IOCD
083F0233	INIT LOAD	NS Init Load

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
083F0234	LOAD STAT	NS Load Status
083F0814	TR_INQUIRY	Translate Inquiry
083F0816	TR_REPLY	Translate Reply
08410210	RNAA	Request Network Address Assignment
08410220	NFY SESEND	Notify Sessend
08410237	LOAD REQRD	NS Load Required
08410240	ADDNR	Add Network Resource
08410243	IPL INIT	NS IPL Init
08410244	IPL TEXT	NS IPL Text
08410245	IPL FINAL	NS IPL Final
08410246	IPL ABORT	NS IPL Abort
08410286	RDELETENR	Request Delete Network Resource
08410287	LOST CTLPT	Lost Control Point
08410289	ROUTE_INOP	Network Services Route Inoperative
0841028A	REQACTCDRM	Request ACTCDRM
08410304	REQMS	Request Maintenance Statistics
08410305	LINKLVL2	Enter Test Mode (LL2)
08410307	REQ RTTEST	Request Route Test
08410311	MS SCV	Maintenance Services Set Control Vector
08410384	RECFMS	Record Formatted Maintenance Statistics
08410385	RECTR	Record Test Results
08410386	ER TESTED	Explicit Route Tested
0841038D	NMVT	Network Manager Vector Transport
08810387	REQ ECHO	Request Echo Test
08810389	ECHO TEST	Echo Test
08810601	CINIT	Control Initiate [®]
08810602	CTERM	Control Terminate
08810620	NOTIFY	Notify
08810629	CLEANUP	Cleanup
08810680	INIT OTHER	Initiate(Other)
08810681	INIT SELF	Initiate(Self) Format 1
08810682	TERM OTHER	Terminate(Other)
08810683	TERM SELF	Terminate(Self)
08810685	BIND FAIL	Bind Failure
08810686	SESS START	Session Started
08810687	UBIND FAIL	Unbind Failure
08810688	SESS ENDED	Session Ended
08810810	FORWARD	Forward Request
08810812	DELIVER	Deliver Request
08810814	CNM	CNM Request

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
08812601	BFCINIT	BF Control Initiate
08812629	BFCLEANUP	BF Cleanup
08812681	BFINIT	BF Initiate
08812683	BFTERM	BF Terminate
08812686	BFSESSST	BF Session Start
08812688	BFSESEND	BF Session End
0881268C	BFSESSINFO	BF Session Information Request
08818620	CD NOTIFY	Cross-Domain Notify
08818627	CD DSRLST	Cross-Domain Direct Search List
08818640	CDINIT OTH	Cross-Domain Initiate (Other)
08818641	CDINIT	Cross-Domain Initiate
08818643	CDTERM	Cross-Domain Terminate
08818645	CDSSF	Cross-Domain Session Setup Failure
08818646	CDESSST	Cross-Domain Session Started
08818647	CDSTF	Cross-Domain Session Takedown Failure
08818648	CDESESEND	Cross-Domain Session Ended
08818649	CDTAKEDOWN	Cross-Domain Takedown
0881864A	CDTD COMP	Cross-Domain Takedown Complete
0881864B	CDCINIT	Cross-Domain Control Initiate
09050000	NCLSA	Network Control Lost Subarea
09060000	ER INOP	Explicit Route Inoperative
09060000	ANS	Auto Network Shutdown Started
09070000	ANSC	Auto Network Shutdown Complete
09080000	LOST PATH	Lost Path
09090000	ER TEST	Explicit Route Test
090A0000	ER TST RPY	Explicit Route Test Reply
090B0000	ER ACT	Explicit Route Activate
090C0000	ER ACT RPY	Explicit Route Activate Reply
090D0000	ACTVR	Activate Virtual Route
090E0000	DACTVR	Deactivate Virtual Route
090F0000	ER OP	Explicit Route Operative
09510000	SW TO NCP	Switch Line to NCP Mode
09520000	SW TO EP	Switch Line to EP Mode
0A040000	LUSTAT	LU Status
0A050000	RTR	Ready to Receive
0A700000	BIS	Bracket Initiation Stopped
0A710000	SBI	Stop Bracket Initiation
0A800000	QEC	Quiesce at End of Chain
0A810000	QC	Quiesce Complete
0A820000	RELQ	Release Quiesce

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
0A830000	CANCEL	Cancel
0A840000	CHASE	Chase
0AC00000	SHUTDOWN	Shutdown
0AC10000	SHUTC	Shutdown Complete
0AC20000	RSHUTD	Request Shutdown
0AC80000	BID	Bid
0AC90000	SIGNAL	Signal
0B0D0000	ACTLU	Activate LU
0B0E0000	DACTLU	Deactivate LU
0B110000	ACTPU	Activate PU
0B120000	DACTPU	Deactivate PU
0B140000	ACTCDRM	Activate CDRM
0B150000	DACTCDRM	Deactivate CDRM
0B310000	BIND	Bind Session
0B320000	UNBIND	Unbind Session
0BA00000	SDT	Start Data Traffic
0BA10000	CLEAR	Clear Session
0BA20000	STSN	Set and Test Sequence Numbers
0BA30000	RQR	Request Recovery
0BC00000	CRV	Cryptography Verify
0C0102A0	AM ALLORSC	Allocate Resource
0C0102A1	AM FREERSC	Free Resource
0C0102A2	AM SETRT	Set Routable State
0C0102A3	AM RESETRT	Reset Routable State
0C0102A4	AM SC EXIT	Configuration Services Exit AMRU
0C0102A5	DDDLU RU	Secondary LU exit AMRU
0C010480	RECMD	Record Measurement Data
0C410201	AM CS	Config SVCS
0C410206	AM GAINGWN	Gained GWN
0C410207	AM LOSTGWN	Lost GWN
0C410208	AM DEACTXF	Deactivate Transforms
0C410210	AM RNAA	Request Network Address Assignment
0C410212	AM CONNECT	Connect
0C410213	AM DISCNCT	Disconnect
0C410214	AM INIT_PU	Initiate PU
0C410266	XID	AM Exchange ID
0C410268	AM XCF	AM XCF
0C4102BD	AM ADDLINK	Add Link
0C4102BE	AM ADDLSTA	Add Link Station
0C4102BF	DELETENR	Delete Network Resource

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
0C4102CD	AM REQDUMP	Request Dump
0C4102CE	AM CONDL0D	Request Conditional Load
0C4102CF	AM UNCDL0D	Request Unconditional Load
0C410601	AM OPNACB	Open ACB
0C410602	AM CLSACB	Close ACB
0C800700	AM VCNS LREQ	VCNS Logon Request
0C800701	AM VCNS LRSP	VCNS Logon Response
0C810619	AM ADRQCMP	Address Request Complete
0C810620	AM RESUME	Resume
0C810629	AM CLEANUP	Cleanup
0C810643	AM GENTERM	Termination Placeholder
0C810680	AM REALLOC	Reallocate
0C810681	SETUP	Generic Session Initiation
0C810801	AM NOTIFY	Notify
0C810A00	API SETLST	SETLOGON(START)
0C810A01	API SETLSP	SETLOGON(STOP)
0C810A02	API SETLQS	SETLOGON(QUIESCE)
0C810A03	API SETPER	SETLOGON(PERSIST)
0C810A04	API SETNPE	SETLOGON(NPERSIST)
0C810A05	API SETGNA	API SETLOGON(GNAMEADD)
0C810A06	API SETGND	API SETLOGON(GNAMEDEL)
0C810A10	API SIMLOG	SIMLOGON
0C810A20	API OPNACQ	OPNDST(ACQUIRE)
0C810A21	API OPNACC	OPNDST(ACCEPT)
0C810A22	API OPNRES	OPNDST(RESTORE)
0C810A30	API INQLOG	INQUIRE(LOGONMSG)
0C810A31	API INQDVC	INQUIRE(DEVCHAR)
0C810A32	API INQCNT	INQUIRE(COUNTS)
0C810A33	API INQTOP	INQUIRE(TOPLOGON)
0C810A34	API INQCID	INQUIRE(CIDXLATE)
0C810A35	API INQTRM	INQUIRE(TERMS)
0C810A36	API INQAPS	INQUIRE(APPSTAT)
0C810A37	API INQSPM	INQUIRE(SESSPARMS)
0C810A38	API INQSKY	INQUIRE(SESSKEY)
0C810A39	API INQDPY	INQUIRE(DISPLAY)
0C810A3A	API INQPER	INQUIRE(PERSESS)
0C810A3B	API INQNQN	INQUIRE(NQN)
0C810A3C	API INQSNM	INQUIRE(SESSNAME)
0C810A3D	API INQSTA	INQUIRE(STATUS)
0C810A40	API INTERP	INTRPRET

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
0C810A50	API CLSPAS	CLSDST(PASS)
0C810A51	API CLSRLS	CLSDST(RELEASE)
0C810A60	API SESONC	SESSIONC
0C810A70	API SNDCMD	SENDCMD
0C810A75	API SEND	SEND
0C810A80	API RCVCMD	RCVCMD
0C810A85	API RECEIV	RECEIVE
0C810A90	API REQSES	REQSESS
0C810AA0	API OPNSEC	OPNSEC
0C810AB0	API TRMSES	TERMSESS
0C810AC0	API RSETSR	RESETSR
0C810AD0	API CHGEAF	CHANGE (ENDAFFIN)
0C810AD1	API CHGEF	CHANGE (ENDAFFNF)
0D010000	AM VR INOP	Virtual Route Inoperative
0D0B0000	AM REQ ERA	Request Explicit Route Activate
0D0E0000	AM REQ VRD	Request Virtual Route Deactivate
0DFF0000	AM VR STAT	Virtual Route Status
0E010000	AM PCE	Purge Chain Element
0F010000	AM NFY SLT	Notify (Schedule LOSTERM Exit)
0F020000	AM SSA	Set Session Address
0F030000	AM SSADISC	Set Session Address and Disconnect
0F040000	AM OSA	Override Session Address
0F050000	AM PWQ	Purge Wait Queue
0F060000	AM FLUSH	Flush Virtual Route
0F310000	AM GBIND	Generic BIND
0F320000	AM GUNBIND	Generic UNBIND
10400004	AGT SRHCMP	Subarea search for resource complete
FCC1C3D9	IPS ACR	CDINIT RESP AMRU
FCC1C3E2	IPS ACS	ACT_CP_CP_SESS_V
FCC1D3E2	IPS ALS	ALERT_SIGNAL_V
FCC1E2D9	IPS ASR	ACT_CP_SVR_SESS_V
FCC2C3E2	IPS BCS	BEGIN_CP_STATUS_V
FCC2D9C9	IPS BRI	BROADCAST IPS
FCC2D5D7	IPS BNP	BN_SESS_RPY
FCC2D5D8	IPS BNQ	BN_SESS_REQ
FCC3C2D5	IPS CBN	CACHE_BN_INFO
FCC3C3E6	IPS CCW	CONTINUE_CW_V
FCC3C4E8	IPS CDY	CDRSC_DISPLAY
FCC3C8C6	IPS CHF	CHAIN_FLOW_V
FCC3C8D2	IPS CHK	CHKPT_START_V

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
FCC3C8D9	IPS CHR	CHAIN_FLOW_RPY_V
FCC3D3C5	IPS CLE	CLEANUP_V
FCC3D4C1	IPS CMA	PROCESS_COSMAP_V
FCC3D6E2	IPS COS	DEFINE COS
FCC3E2C1	IPS CSA	CP_CP_SESS_ACT_V
FCC3E2C4	IPS CSD	CP_CP_SESS_DEACT_V
FCC3E2C8	IPS CSH	CACHE_SEARCH
FCC3E2D9	IPS CSR	CACHE_SEARCH_RPY
FCC3E2F0	IPS CS0	CSS_DISCR_INIT_REQ_V
FCC3E2F1	IPS CS1	CSS_DISCR_INIT_RPY_V
FCC3E2F2	IPS CS2	CSS_TOPO_INIT_REQ_V
FCC3E2F3	IPS CS3	CSS_TOPO_INIT_RPY_V
FCC3E2F4	IPS CS4	CSS_IO_REQ_V
FCC3E2F5	IPS CS5	CSS_IO_RPY_V
FCC3E2F6	IPS CS6	CSS_TERM_V
FCC3E2F7	IPS CS7	CSS_API_DATA_V
FCC3E2F8	IPS CS8	CSS_INTER_STACK_DATA_V
FCC3E2F9	IPS CS9	CSS_ABEND_START_V
FCC3E2C1	IPS CSA	CSS_ABEND_V
FCC3E2C2	IPS CSB	CSS_MDS_DATA_V
FCC3E2C3	IPS CSC	CSS_MST_INIT_REQ_V
FCC3E2C5	IPS CSE	CSS_MODIFY_OSIEVENT_V
FCC3E2C6	IPS CSF	CSS_STOP_ONGOING_V
FCC3E3C6	IPS CTF	CP_CP_TP_FAILURE_V
FCC4C1D5	IPS DAN	DISPLAY_AJNLT_V
FCC4C1D7	IPS DAP	DAP_TP_WORK_V
FCC4C1E4	IPS DAU	DEALLOCATE_ABEND_USER_V
FCC4C3D6	IPS DCO	DISPLAY_APPNCOS_V
FCC4C3E2	IPS DCS	DEACTIVATE_CP_CP_SESS_V
FCC4C3E6	IPS DCW	DRIVE_CONWINNER_V
FCC4C6C1	IPS DFA	DEFINE_ADJCLUST_V
FCC4C9C1	IPS DIA	DISPLAY_ADJCLUST_V
FCC4D3D7	IPS DLP	DLUR_PATH_SWITCH_COMP
FCC4D3D9	IPS DLR	DLUR_STATUS
FCC4D4C2	IPS DMB	DSME_BN_SELECT_V
FCC4D4C3	IPS DMC	DSME_CDS_SELECT_V
FCC4D4C9	IPS DMI	DSME_ICN_SELECT_V
FCC4D8C5	IPS DQE	DEQUEUE
FCC4D9E5	IPS DRV	DATA_RECOVERED_V
FCC4E2C9	IPS DSI	DISPLAY_SRCHINFO

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
FCC4E2D4	IPS DSM	DSME_IAUTH_V
FCC4E2D9	IPS DSR	DEACT_CP_SVR_SESS_V
FCC5D9D7	IPS ERP	RTP Exprec Purge IPS
FCC6E2C4	IPS FSD	FLUSH_CP_SVR_DATA_V
FCC6E2E3	IPS FST	RTP Free Storage IPS
FCC7C3D9	IPS GCR	GENERIC CACHE SEARCH REPLY
FCC7C3E2	IPS GCS	GENERIC CACHE SEARCH
FCC7E4E2	IPS GUS	GENERIC CACHE SEARCH UPDATE REQUEST
FCC8C3D9	IPS HCR	RTP_CPNAME Change IPS
FCC9D5D6	IPS INO	INIT_OTHER
FCC9D6C3	IPS IOC	INIT_OTHER_COMP
FCD4D7C9	IPS MPI	MNPS_Pipe_Info IPS
FCD4E2C9	IPS MSI	MNPS_SESSINFO
FCD5E4E3	IPS NUT	NLP with Unknown TCID value
FCD5E5D9	IPS NVR	NON_VERIFY_REQ
FCD6E3C3	IPSOTC	ORDERLY_SESSIONS_TERM_C
FCD6E3C4	IPS OTD	OUTPUT TDU SIGNAL
FCD6E3D9	IPS OTR	ORDERLY_SESSIONS_TERM_R
FCD7C3D8	IPS PCQ	PCID_QUERY
FCD7C3D9	IPS PCR	PCID_QUERY_REPLY
FCD7D3E3	IPS PLT	PURGE_LOCATE_TIMER
FCD7D5C4	IPS PND	PEND_SC
FCD7D5D3	IPS PNL	PROCESS_NNSLIST
FCD7D9C3	IPS PRC	PROCEED
FCD7D9E5	IPS PRV	PROV_SC
FCD8C5C4	IPS QED	QUEUED
FCD9C1C3	IPS RAC	RESOURCE_AVAILABLE_COMP
FCD9C1E5	IPS RAV	RESOURCE_AVAILABLE
FCD9C3E2	IPS RCS	RTP_Connection Setup
FCD9C9E5	IPS RIV	RTP_Inactivation Request/Reply IPS
FCD9D3C6	IPS RLF	REQ_LAST_FRSN_V
FCD9D4D9	IPS RMR	REQUEST_MULTIPLE_ROUTES
FCD9D5C3	IPS RNC	Req_NonPersistent CLOSE IPS Mapping
FCD9D9C3	IPS RRC	RELEASE_REQUEST_COMP
FCD9D9D8	IPS RRQ	RELEASE_REQUEST
FCD9E2D5	IPS RSN	RSN_NOTIFY
FCD9E2D9	IPS RSR	RTP Route_Setup Req for PS
FCD9E2E3	IPS RST	RTP_Setup IPS
FCD9E2E4	IPS RSU	Route_Setup IPS

Table 45. Control point/control block operation codes (CPCBOPC) (continued)

CPCBOPC	Message display	Function
FCD9E3C9	IPS RTI	RTP_TG Inoperative IPS
FCE2C1C3	IPS SAC	SUBAREA CLEANUP
FCE2C3D9	IPS SCR	SEARCH_REQ
FCE2C4E8	IPS SDY	DISPLAY_SESSIONS
FCE2C8D9	IPS SHR	REQUEST SINGLE HOP ROUTE
FCE2C9D9	IPS SIR	SESS_INIT_INFO_REQ
FCE2D9C3	IPS SRC	SESSION_REQUEST_COMP
FCE2D9D8	IPS SRQ	SESSION_REQUEST
FCE2D9E2	IPS SRS	CP_SVR_SESS_STAT_V
FCE4D4D9	IPS UMR	UPDATE MODES
FDC1C3C3	IPS ACC	ADJACENT_CP_CONTACTED
FDC1C3E4	IPS ACU	ADJACENT_CP_UPDATED
FDC3D7E2	IPS CPS	CP_STATUS
FDC3D9D8	IPS CRQ	CRR_REQUEST
FDC4E2D5	IPS DSN	DIRECTORY SERVER NOTIFY
FDD9C1D3	RAL	RTP_Allocation Request/Reply IPS
FDD9C1E3	RAT	RTP_Attach Request/Reply IPS
FDD9C4D3	RDL	RTP_Deallocation Request/Reply IPS
FDD9C4E3	IPS RDT	RTP_DEATTACH
FDD9C9D6	RIO	RTP_Inoperative IPS
FDD9E2C3	RSC	RSCV Scan IPS
FDD9E2D9	IPS RSR	RES_REGISTRATION
FDE2D9D6	IPS SRO	CPSVRMGR SESS Outage
FDE2D9D7	IPS SRP	CP_SVR_PROT_VIOL_V
FDE2D9E3	IPS SRT	TDU Error from DLUR
FDE2E3D7	IPS STP	START_TP
FDE2E3D9	IPS STR	START_TP_REPLY
FDE3C7E4	IPS TGU	TG UPDATE
FDE3D7C6	IPS TPE	TP_ENDED
FDE3D7D5	IPS TPN	TP_NOTIFY
FDE4D7C4	IPS UPD	UPDATE_DIRECTORY
FF000000	VECTOR	VECTOR FMD Request Units

Appendix E. Storage and control block ID codes

This appendix lists the control block ID codes of the VTAM program.

VTAM control block ID codes

You can identify certain VTAM control block types in a storage dump by examining an identification code in the first byte of the control block (offset 0). The control block identification codes are shown in the following table. If codes are duplicate, use other means (such as the operating environment or the control block's context) to determine the type of control block.

Note: Internal codes that are used only by the product support organization to assist in internal flow diagnosis are not included in Table 46.

Table 46. Control block ID codes

ID	Control block
00	SONCB
00	RPL
01	RPH
03	FMCB
05	VRBLK
06	ICNCB
07	LDNCB
09	HALCB
0A	BSCLB
0B	VLNCB
0C	PCLCB
0D	PRWCB
0E	TRGCB
0F	ACDEB
10	UECB
11	DYPAB
13	TRAC
14	ERT
15	ISTAUNCB
16	ISTIPNCB
17	ISTAUCPL
19	PDVT
1A	CHAIN
1B	TGCB
1E	RCE
24	OCW

Table 46. Control block ID codes (continued)

ID	Control block
26	PICB
29	LMPCB
2B	RAQ
2C	PAQ
2D	SAT
2E	AHNCB
2F	ISTALNCB
41	PUSCB
43	PLSCB
45	POIA
46	POCB
47	POMCB
48	PORCB
49	POWE
4C	ERCT
4D	TGE
4E	VRWSE
50	DLRPL
52	LUCB
54	RUPE
54	TUNB
58	TQE
5A	PRQAB
5B	PRBLK
5C	CPCB
5E	IEF
60	NCSPL
61	PST
62	AMU
62	NSSCB
62	NSICB
63	SMP
64	OCB
65	NACP
66	CAB
67	CANT
68	RSQE
69	RDTPPL
6A	NOSPL
6B	SSIB

Table 46. Control block ID codes (continued)

ID	Control block
6B	CAR
6C	IOSIB
6D	ASRIT
6D	RANT
6E	GWIT
6F	RARB
73	XCB
74	PWK
75	WRE
77	ADJSR
78	ADJSS
7B	RIB
81	EXLST
96	SIBXN
97	SIBIX
98	SIB
99	TSCB
9A	TSPL
9B	LSCB
9C	CNCB
A0	ACB
A2	RNCA
A3	ALCA
BD	UDT
BE	INT1
BF	COS
C0	LOGMD
C1	RPL6X
CA	GRPCB
D0	NIB
EB	AUTOE
FA	RPNCB
FC	TLNCB
FD	RWNCB
FE	XCNCB
FF	OCA
ARB	ARB
ART	ART
AUAU	AULIN_ARRAY
AULN	AULIN

Table 46. Control block ID codes (continued)

ID	Control block
AUVT	AUVT
CLK	CLK
CLWB	HCLW
DAPT	DAPTR
FLU	FLU
FQPT	FQPTB
FRSR	FRSRC
FRTP	FRTP
HIT	HIT
HTMB	HTMBK
LNKT	LNKTB
LSP	LSPL
LUTB	LUTAB
MRPF	MRPFA
PRT	PRTCB
None	RCM
RCOR	RCORS
RUR	RUR
RTPT	RTPTB
SAP	SAPCB
SGMN	SGMNT
SND	SND
TIMB	TIMBK

Appendix F. Installing dump analysis and VIT analysis tools

The dump analysis and VIT analysis tools are used for diagnosing software failures.

Concatenating target data sets used in the installation

Table 47 shows the target data sets that contain the data necessary to set up the z/OS Communications Server dump analysis and the VIT analysis tool. You need to concatenate the target data sets into the DDNAME statements shown.

Table 47. Target data sets for dump and trace tools

Target data set	Action	DDNAME	Comment
SYS1.SBLSTBL0	Concatenate	ISPTLIB	Contains compiled tables, keylists, and commands
SYS1.SBLSCLI0	Concatenate	SYSPROC	Contains CLISTs and REXX execs
SYS1.SBLSPNL0	Concatenate	ISPPLIB	Contains compiled panels
SYS1.SBLSMSG0	Concatenate	ISPMLIB	Contains compiled ISPF messages

Use a LOGON PROC to concatenate the data sets. If you create a new LOGON PROC, you need to log off and then log back on for the PROC to take effect.

See z/OS MVS IPCS Customization for an example of a LOGON PROC.

Customizing IPCS interface

If you want a customized interface to be active to select the z/OS Communications Server dump analysis commands, customize the IPCS panel BLSPPRIM by adding the highlighted lines in Figure 269 on page 642 to create and activate option 7 on the IPCS Primary Option Menu as shown in Figure 270 on page 643. This modification allows you to access VTAMMAP directly for dump processing. When this option is selected, control is passed to the ISTDE01 EXEC. This EXEC controls the IPCS panels for the dump formatter.

For information regarding TCP/IP IPCS CLISTs, see z/OS Communications Server: IP Diagnosis Guide.

Note: This sample is not necessarily identical to the one on your system.

```

)ATTR
/* ===== */
/* 5685-001 This panel is "Restricted materials of IBM" */
/* (C) Copyright IBM Corporation 1988 */
/* Licensed materials - property of IBM */
/* Refer to copyright instructions, form number G120-2083 */
/* ===== */
† TYPE(INPUT) INTENS(HIGH) CAPS(OFF) JUST(LEFT) PAD(NULLS)
@ TYPE(TEXT) COLOR(GREEN) INTENS(LOW)
)BODY
%----- IPCS PRIMARY OPTION MENU -----
%OPTION ==>†ZCMD

%*****
% 0 +DEFAULTS - Specify default dump and options @* USERID - &ZUSER
% 1 +BROWSE - Browse dump data set @* DATE - &ZDATE
% 2 +ANALYSIS - Analyze dump contents @* JULIAN - &ZJDATE
% 3 +SUBMIT - Submit problem analysis job to batch @* TIME - &ZTIME
% 4 +COMMAND - Enter IPCS subcommand or CLIST @* PREFIX - &ZPREFIX
% 5 +UTILITY - Perform utility functions @* TERMINAL- &ZTERM
% 6 +DUMPS - Manage dump inventory @* PF KEYS - &ZKEYS
% 7 +VTAM - VTAM dump analysis commands %*****
% T +TUTORIAL - Learn how to use the IPCS dialog
% X +EXIT - Terminate using log and list defaults

+Enter%END+command to terminate IPCS dialog
)INIT
&ZPRIM = YES /* Always a primary option menu */
&ZHTOP = BLSPhelp /* Tutorial table of contents */
.CURSOR = ZCMD
.HELP = BLSPhelp
&ZHINDEX = &Z /* No tutorial index is supplied */
)PROC
&PASSLIB = &Z
IF (&ZBCS = YES, NO)
&PASSLIB = PASSLIB
&ZSEL = TRANS( TRUNC (&ZCMD, '.')
0, 'PGM(BLSGSCMD) PARM(%BLSCSETD)'
1, 'PGM(BLSLDISP) NEWAPPL(BLSL) &PASSLIB'
2, 'PANEL(BLSPSCRN)'
3, 'PANEL(BLSPBKGD)'
4, 'PANEL(BLSPDSLE)'
5, 'PANEL(BLSPUTIL)' /* %00A*/
6, 'PGM(BLSGDUIIN)'
7, 'PGM(BLSGSCMD) PARM(%ISTDE01) NEWAPPL(ISTD) &PASSLIB'
t, 'PGM(ISPTUTOR) PARM(BLSPTUTR)'
T, 'PGM(ISPTUTOR) PARM(BLSPTUTR)'
, , ,
x, 'EXIT'
X, 'EXIT'
*, '?' )
)END

```

Figure 269. Sample IPCS panel BLSPPRIM customization

```

-----IPCS PRIMARY OPTION MENU-----
OPTION ==> _

0  DEFAULTS   - Specify default dump and options
1  BROWSE     - Browse dump data set
2  ANALYSIS   - Analyze dump contents
3  SUBMIT     - Submit problem analysis job to batch
4  COMMAND    - Enter IPCS subcommand or CLIST
5  UTILITY    - Perform utility functions
6  DUMPS     - Manage dump inventory
7  VTAM      - VTAM dump analysis commands
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 270. Addition of option 7 to the IPCS primary option menu

Verifying dump formatter panels

To verify that dump formatter panels are set up correctly, choose option 7 on the menu shown in Figure 270.

The first ISPF panel you should see is shown in Figure 271.

```

                          VTAMMAP Analysis Menu

Select one of the following items. Then press Enter.

— 1. APPC . . - APPLCONV, PARTNRLU, APPLMODE, APPMODAL
   2. APPN . . - APPNBASE, FNDADJCP, FNDANDCB, FNDCOS, FNDDECB, etc
   3. General. - HOST, VTAM, VTBASIC, VTFNDMOD, VTMODS, VITAL, etc
   4. Queues . - PABSCAN, VTCVTPAB, VTREADYQ
   5. Resource - RDTCHECK, RDTFULL, RDTHIER, RDTSUM, VTNODE
   6. Session. - ATMDATA, FINDDSIB, FINDSIB, MNPS, SES, SIBCHECK
   7. Search . - SRTFIND
   8. Storage. - SPANC, STORAGE, VTBUF, VTRPH
   9. CSM . . - CSMALL, CSMBUF, CSMCMPID, CSMOWNER, CSMPOOL
  10. Waits. . - VTWRE
  11. ERs/VRs. - ROUTES, VTVRBLK
  12. CLISTs. .- ISTVABND, ISTVDUMP, ISTVMAP, ISTVSAVE, ISTVSLIP
  13. APPN2. .- TRSTRACE

(C) Copyright IBM Corporation 1993,2006. All rights reserved.
Command ==> _____
F1=Help  F2=Split  F3=Exit  F9=Swap  F12=Cancel

```

Figure 271. Main menu for selecting dump options

Press the PF1 key to verify that the appropriate help panel is displayed.

Customizing ISPF interface

If you want a customized interface to be active to select the z/OS Communications Server trace analysis commands, customize the ISPF panel ISR@PRIM by adding the highlighted lines shown in Figure 272 on page 645 to create and activate option V on the ISPF/PDF Primary Option Menu as shown in Figure 273 on page 646. When this option is selected, control is passed to the ISTTE01 EXEC. This EXEC controls the ISPF panels for trace formatter.

Note: The samples shown in Figure 272 on page 645 and Figure 273 on page 646 are not necessarily identical to the ones on your system.


```

)ATTR
+ TYPE(TEXT) COLOR(GREEN) INTENS(LOW)
)BODY
%----- SAMPLE ISPF/PDF PRIMARY OPTION MENU -----
%OPTION ==>_ZCMD
%
%
%                                +USERID  - &ZUSER
% 0 +ISPF PARMS  - Specify terminal and user parameters +TIME    - &ZTIME
% 1 +BROWSE     - Display source data or output listings +TERMINAL - &ZTERM
% 2 +EDIT       - Create or change source data          +PF KEYS  - &ZKEYS
% 3 +UTILITIES  - Perform utility functions
% 4 +FOREGROUND - Invoke language processors in foreground
% 5 +BATCH      - Submit job for language processing
% 6 +COMMAND    - Enter TSO command or CLIST
% 7 +DIALOG TEST - Perform dialog testing
% 8 +LM UTILITIES- Perform library administrator utility functions
% 9 +IBM PRODUCTS- Same as option S (SER PRODUCTS)
% 10 +SCLM      - Software Configuration and Library Manager
% C +CHANGES   - Display summary of changes for this release
% V +VTAM       - VTAM trace analysis commands
% T +TUTORIAL   - Display information about ISPF/PDF
% S +SER PRODUCTS- Southeast Region product options
% I +SER IC TOOLS- Southeast Region Info-Center and Toolkits
% P +RPM        - Regional Problem Management
% X +EXIT       - Terminate ISPF using log and list defaults
%
+Enter%END+command to terminate ISPF.
%
)INIT
.HELP = ISR00003
&ZPRIM = YES      /* ALWAYS A PRIMARY OPTION MENU */
&ZHTOP = ISR00003 /* TUTORIAL TABLE OF CONTENTS */
&ZHINDEX = ISR91000 /* TUTORIAL INDEX - 1ST PAGE */
&ZSCLMPRJ = &Z
VPUT (ZHTOP,ZHINDEX,ZSCLMPRJ) PROFILE
)PROC
&ZQ = &Z
IF (&ZCMD ^= ' ')
  &ZQ = TRUNC(&ZCMD, '.')
  IF (&ZQ = ' ')
    .MSG = ISRU000
&ZSEL = TRANS( &ZQ
  0, 'PANEL(ISPOPTA)'
  1, 'PGM(ISRBRO) PARM(ISRBRO01)'
  2, 'PGM(ISREDIT) PARM(P,ISREDM01)'
  3, 'PANEL(ISRUTIL)'
  4, 'PANEL(ISRFPA)'
  5, 'PGM(ISRJB1) PARM(ISRJPA) NOCHECK'
  6, 'PGM(ISRPTC)'
  7, 'PGM(ISPYXDR) PARM(ISR) NOCHECK'
  8, 'PANEL(ISRLPRIM)'
  9, 'PANEL(SERPP000)' /* CHANGED HERE? FROM ISRDIIIS */
  10, 'PGM(ISRSCLM) NOCHECK'
  C, 'PGM(ISPTUTOR) PARM(ISR00005)'
  V, 'CMD(%ISTTE01) NEWAPPL(ISTT) &PASSLIB'
  T, 'PGM(ISPTUTOR) PARM(ISR00000)'
  S, 'PANEL(SERPP000)'
  I, 'PANEL(SERIC000)'
  P, 'CMD(%SRRPM)'
  , , ,
  X, 'EXIT'
  *, '?' )
&ZTRAIL = .TRAIL
)END

```

Figure 272. Sample ISPF panel ISR@PRIM customization

Note: This sample is not necessarily identical to the one on your system.

```
----- ISPF/PDF PRIMARY OPTION MENU -----
OPTION ==>
      0 ISPF PARMS - Specify terminal and user parameters   USERID - USERID
      1 BROWSE    - Display source data or output listings  TIME    - 9:29
      2 EDIT      - Create or change source data           TERMINAL - 3278
      3 UTILITIES - Perform utility functions              PF KEYS - 12
      4 FOREGROUND - Invoke language processors in foreground
      5 BATCH     - Submit job for language processing
      6 COMMAND   - Enter TSO command or CLIST
      7 DIALOG TEST - Perform dialog testing
      8 LM UTILITIES- Perform library administrator utility functions
      9 IBM PRODUCTS- Same as option S (SER PRODUCTS)
     10 SCLM     - Software Configuration and Library Manager
      C CHANGES  - Display summary of changes for this release
      V VTAM     - VTAM trace analysis commands
      T TUTORIAL - Display information about ISPF/PDF
      S SER PRODUCTS- Southeast Region product options
      I SER IC TOOLS- Southeast Region Info-Center and Toolkits
      P RPM      - Regional Problem Management
      X EXIT     - Terminate ISPF using log and list defaults

Enter END command to terminate ISPF.
```

Figure 273. Addition of option V to the ISPF/PDF primary option menu

Verifying trace formatter panels

To verify that trace formatter panels are set up correctly, choose option V on the menu shown in Figure 273.

The first ISPF panel you should see is shown in Figure 274.

```
VTAM Internal Trace Analysis

Select one of the following items. Then press Enter.

— 1. Storage Analysis
   2. Request/response unit counting
   3. VIT extraction
   4. Input complete

(C) Copyright IBM Corporation 1993,2002. All rights reserved.
Command ==>
F1=Help    F2=Split    F3=Exit    F9=Swap    F11=Retrieve F12=Cancel
```

Figure 274. Main menu for selecting trace parameters

Press the PF1 key to verify that the appropriate help panel is displayed.

Note: It is recommended that you position the command line at the bottom of the screen using ISPF PARMs option DISPLAY and changing the 'COMMAND LINE PLACEMENT ==> ASIS' to BOTTOM to improve readability.

Appendix G. Problem topics in other libraries

Table 48. Related information on problem topics in other libraries

Topic	See
3174 Controller	<i>3174 Functional Description</i>
Abend codes	<i>z/OS MVS System Codes</i>
Abend dump	<i>z/OS MVS JCL Reference</i>
Alerts	<i>Tivoli NetView for z/OS Version 5.2 Command Reference Volumes 1 & 2</i> <i>SNA Network Product Formats</i>
CCW trace	<i>z/OS MVS Diagnosis: Tools and Service Aids</i>
Channel programs	<i>IBM 4361 Processor Communication Adapter</i> <i>IBM 9370 Information System: Telecommunications Subsystem Description and Reference</i> <i>Principles of Operation</i> manual for your communication controller <i>Principles of Operation</i> manual for your operating system
CNOS return codes	<i>z/OS Communications Server: SNA Programmer's LU 6.2 Guide</i>
Communication scanner output	<i>NCP, SSP, and EP Diagnosis</i> <i>NCP Reference Summary and Data Areas</i> <i>Principles of Operation</i> manual for your communication controller
Coupling facility structures dump	<i>z/OS MVS System Commands</i>
Directory services management exit	<i>z/OS Communications Server: SNA Customization</i>
Dump collecting, formatting, and printing	<i>z/OS MVS Diagnosis: Tools and Service Aids</i>
ERP	<i>SYS1.LOGREC Error Recording</i>
Exception request (EXR)	<i>SNA Network Product Formats</i>
First Failure Support Technology (FFST)	<i>First Failure Support Technology for VM and MVS Operator's Guide</i>
Full-screen mode	<i>z/OS TSO/E Programming Services</i>
Generalized trace facility (GTF)	<i>z/OS MVS Diagnosis: Tools and Service Aids</i>
Generic alerts	See <i>Alerts</i> . <i>NCP, SSP, and EP Diagnosis Guide</i>
Hung NCP Hung resources attached to an NCP, Hung sessions	
IEBGENER utility	<i>MVS Utilities</i>
Intensive mode recording	<i>NCP, SSP, and EP Diagnosis Guide</i>
I/O control blocks	This information is available in the MVS data areas documentation, which is available at the following website: z/OS Internet Library .
I/O traces	This information is available in the MVS data areas documentation, which is available at the following website: z/OS Internet Library .

Table 48. Related information on problem topics in other libraries (continued)

Topic	See
IPCS, running in batch mode, IPCS commands	z/OS MVS IPCS Commands z/OS MVS IPCS User's Guide
IPCSPRNT	z/OS MVS IPCS Customization
Job control language (JCL)	z/OS MVS JCL Reference z/OS MVS JCL User's Guide
Line trace records	<i>NCP, SSP, and EP Diagnosis Guide</i>
LIST Service Aid	z/OS MVS Diagnosis: Tools and Service Aids
LOGDATA option	z/OS MVS IPCS Commands
LOGREC	<i>SYS1.LOGREC Error Recording</i>
MVS IKJxxxx system messages	z/OS MVS System Messages, Vol 9 (IGF-IWM)
MVS system codes	z/OS MVS System Codes
NCP data areas, registers, and codes	<i>NCP and EP Reference Summary and Data Areas, Volumes I and II</i>
NCP dumps, NCP service aids	<i>NCP, SSP, and EP Diagnosis Guide</i>
NCP, tuning	<i>NTuneMON User's Guide</i>
NetView Session Monitor	<i>Tivoli NetView for z/OS Version 5.2 Command Reference Volumes 1 & 2</i>
NMVT RUs	<i>SNA Network Product Formats NetView Operation</i>
Program status words (PSWs)	<i>Principles of Operation</i> manuals
RECMS RU formats	<i>NCP and EP Reference Summary and Data Areas, Volumes I and II</i>
Reshow processing	z/OS TSO/E Programming Services
RTCT	This information is available in the MVS data areas documentation, which is available at the following website: z/OS Internet Library.
RU opcodes	<i>SNA Formats</i>
SDATA options	z/OS MVS System Commands
Session monitor	See <i>NetView Session Monitor</i> .
SLIP dump	z/OS MVS Diagnosis: Tools and Service Aids
SMP	<i>System Modification Program Extended User's Guide</i>
SNA sense codes	<i>SNA Formats</i>
Socket API calls	z/OS Communications Server: IP Programmer's Guide and Reference
Stand-alone dump	z/OS MVS Diagnosis: Tools and Service Aids
STATMON	See <i>NetView Session Monitor</i> .
SVC dump	z/OS MVS System Commands
SVC 93 and SVC 94 entries	z/OS MVS Diagnosis: Tools and Service Aids
Task control blocks	This information is available in the MVS data areas documentation, which is available at the following website: z/OS Internet Library.
TCB, map of	This information is available in the MVS data areas documentation, which is available at the following website: z/OS Internet Library.
TGET options, TGET return codes	z/OS TSO/E Programming Services
TGET/TPUT option flags	z/OS MVS Diagnosis: Reference

Table 48. Related information on problem topics in other libraries (continued)

Topic	See
TPUT options, editing done by	z/OS TSO/E Programming Services

Appendix H. Architectural specifications

This appendix lists documents that provide architectural specifications for the SNA Protocol.

The APPN Implementers' Workshop (AIW) architecture documentation includes the following architectural specifications for SNA APPN and HPR:

- APPN Architecture Reference (SG30-3422-04)
- APPN Branch Extender Architecture Reference Version 1.1
- APPN Dependent LU Requester Architecture Reference Version 1.5
- APPN Extended Border Node Architecture Reference Version 1.0
- APPN High Performance Routing Architecture Reference Version 4.0
- SNA Formats (GA27-3136-20)
- SNA Technical Overview (GC30-3073-04)

For more information, see the AIW documentation page at <http://www.ibm.com/support/docview.wss?rs=852&uid=swg27017843>.

The following RFC also contains SNA architectural specifications:

- RFC 2353 *APPN/HPR in IP Networks APPN Implementers' Workshop Closed Pages Document*

RFCs can be obtained from:

Government Systems, Inc.
Attn: Network Information Center
14200 Park Meadow Drive
Suite 200
Chantilly, VA 22021

Many RFCs are available online. Hardcopies of all RFCs are available from the NIC, either individually or by subscription. Online copies are available using FTP from the NIC at <http://www.rfc-editor.org/rfc.html>.

Use FTP to download the files, using the following format:

```
RFC:RFC-INDEX.TXT  
RFC:RFCnnnn.TXT  
RFC:RFCnnnn.PS
```

where:

- *nnnn* is the RFC number.
- TXT is the text format.
- PS is the postscript format.

You can also request RFCs through electronic mail, from the automated 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.

Appendix I. 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 the z/OS Information 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 the Information Center, as described in "Dotted decimal syntax diagrams."

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users accessing the Information 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

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.

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.

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](http://www.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 xxii.

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. In addition, Info APARs for z/OS documents are in *z/OS and z/OS.e DOC APAR and PTF ++HOLD Documentation*, which can be found at http://publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/Shelves/ZDOCAPAR.

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.

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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Communication network management (CNM) routing table • Logon-interpret routine requirements • Logon manager installation-wide exit routine for the CLU search exit • TSO/SNA installation-wide exit routines • SNA installation-wide exit routines

Writing application programs

Title	Number	Description
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: <ul style="list-style-type: none"> • 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

- 3720 communication controller, alert messages from 401
- 3725 communication controller, alert messages from 401
- 3745 communication controller
 - alert messages from 401
 - bus switching 402

A

- abend (abnormal end)
 - codes
 - 0AB 108, 109
 - 0AC 110
 - 0AD 110
 - 0Ax 60
 - 0C4 8
 - 0Cx 59
 - 15D 110
 - diagnosis procedure 57
 - dump 183
 - finding the abend code 59
 - symptoms 4, 42
 - TSO/VTAM
 - diagnosis procedure 109
 - documentation requirements 109
 - symptoms 105
 - VTAM address space 242
- Abend in user exit ISTECCS, ISTECCSD, ISTECCAA, or ISTECCVR 9
- accessibility 655
- ACF/TAP (trace analysis program) 323
- activation
 - CDRM 18
 - NCP 22
 - network traces 308
- alert messages 401
- ALL VTAMMAP 196
- alternate screen size, PSERVIC coding 118
- analysis and trace analysis tools, installing dump 641
- analyzing dumps with IPCS VTAMMAP 191
- APAR (authorized programming analysis report) 48
- APPLCONV VTAMMAP 197
- application program
 - log 50
 - performance group specification (TSO/VTAM) 120
 - problems 49
- APPLMODE VTAMMAP 198
- APPMODAL VTAMMAP 200
- APPN flows 521
- APPNBASE VTAMMAP 202
- ASCB trace field 325
- ASKQ item, in problem reporting 4
- AUTH operand (TSO/VTAM) 120
- authorized programming analysis report (APAR) 48

B

- backspace key functions improperly 115
- BIND in logon, detecting incorrect portion of 107
- BIND location 118

- BINFMT, coding 107
- bracket in buffer contents trace output 112
- buffer contents trace
 - confidential data 335
 - description 328
 - full 331
 - operation 331
 - output 334
 - partial 331
 - when to use 329
- buffer extents, effect on performance (TSO/VTAM) 120
- buffer pool, analyzing usage 127, 348
- buffer use
 - trace 346, 348
 - trace format, MVS 349
- buffer use display 127
- bus switching 402

C

- CCW trace 92, 310
- change direction indicator in buffer contents trace output 112
- changing screen size in non-full-screen processing 118
- channel program
 - XID
 - X-side 424, 434
 - Y-side 425, 434
- channel programs
 - APPN host-to-host channels 420
 - channel-attached non-SNA 3270 devices 440
 - channel-attached type 2 and 4 PUs 413
 - channel-to-channel adapters 420
 - command words 441
 - HPDT MPC data transfer 438
 - MPC connection, activating 434, 436
 - MPC data transfer 436
 - multipath channels 420
 - PUNS-related 418
- channel-attachment major node, I/O trace 341
- character delete key functions improperly 115
- CICS, logon failure 17
- CLISTS
 - detailed descriptions 196
 - errors, debugging 192
 - help, obtaining online 192
 - ISTVABND 242
 - ISTVDUMP 244
 - ISTVMAP 247
 - ISTVSAVE 249
 - ISTVSLIP 251
 - output, printing 193
 - overview 192
- CLSDST PASS, possible cause for failure 107
- codes, CPCB operation 623
- colons as incorrect output (TSO/VTAM) 116
- commands
 - operator, DISPLAY and MODIFY 123
 - rejected (TSO/VTAM) 105
 - to change screen size in non-full-screen processing (TSO/VTAM) 118

- common problems
 - APPN 22
 - description 4
 - path and routing 157
 - subarea 4
 - VIT analysis tool 44
 - VTAM dump analysis tool 47
- common symptoms 42
- communication scanner processor (CSP)
 - dump 187
 - trace 359
- Communications Server for z/OS, online information xxiii
- component ID 53
- concatenated input 115
- confidential data in buffer contents trace 335
- configuration data set 51
- console log 52
- control blocks
 - formatted in a dump 189
 - ID codes, for VTAM 637
- control operator control block (COPR), finding 197
- control point/control block (CPCB) operation codes 623
- coordinated universal time (Greenwich mean time) 275
- COPR (control operator control block), finding 197
- COS
 - undefined entry, (APPN) 27
 - undefined entry, (subarea) 17
- CP-CP session flows 525
- CPCB operation codes 623
- CPU trace field 325
- CRA 73
- CRALXACT 74
- CRALXPTR fields 74
- cross-domain logon problem 108
- CSP (communication scanner processor)
 - dump 187
 - trace 359
- customer number 55

D

- data
 - display problem (TSO/VTAM) 112
 - incorrect content
 - TSO/VTAM 105
 - VTAM 90
 - incorrect length (TSO/VTAM) 105
 - incorrect translation (TSO/VTAM) 116
 - misplaced, TSO/VTAM 115
- default
 - logmode name in USS command 117
 - screen size 118
- dependent LU server flows 587
- detected waits increase 106
- device workload information, displaying 175
- diagnosing problems in VTAM dump analysis 47
- diagnosing problems, where to begin 3
- directory services flows 533
- disability 655
- dispatching priority, performance group specification 120
- DISPLAY command 128
 - active traces 175
 - buffer pool use 127
 - communications storage manager 128
 - device workload information 175
 - EEDIAG 131
 - Enterprise Extender 128

- DISPLAY command (*continued*)
 - Enterprise Extender connection network unreachable
 - partner information 149
 - HPR route test 150
 - NCP storage 155
 - network ID 156
 - network resource status 156
 - path tables 155
 - pending state 156
 - route status 156
 - route test 163
 - RTPS 150
 - topology database update statistics 167
 - trace status 175
 - TRL 176
 - VTAM storage 175
- DNS, online information xxv
- documentation problem 100
- documenting problems 48, 55
- dump analysis and trace analysis tools, installing 641
- dumps
 - ABDUMP 189
 - analysis of VTAM 191, 192
 - communication scanner processor (CSP) 187
 - dynamic NCP 185
 - enhanced dump analysis tools 191
 - formatted dump procedures 196
 - maintenance and operator subsystem (MOSS) 188
 - MVS
 - abend 183
 - analyzing with VTAMMAP 191
 - formatting and printing 188
 - stand-alone 184
 - SVC 184
 - VTAM control blocks formatted 189
 - network control program (NCP) dump
 - using independent dump utility 187
 - using VTAM dump facility 186
 - static NCP 185
 - tool, dump analysis 191
 - tracing execution sequence of VTAM components in 405
 - dynamic, NCP dump 185

E

- echo test 403
- editing done by TPUT options 116
- EIDs (event IDs) 69
- electronic technical report (ETR) 4
- error recording
 - hardware 403
 - intensive mode 179
 - LOGREC 51
 - NCP 404
 - NMVT alerts 401
- escape character associated with incorrect screen size 119
- event IDs (EIDs) 69
- exception responses 114
- execution sequence of VTAM components in a dump 405

F

- failing module problem, procedure 100
- FINDSIB VTAMMAP 216
- finding the PSW 86
- FINDSIB VTAMMAP 218

flows

- APPN
 - CP-CP session 525
 - dependent LU server 587
 - description 521
 - directory services 533
 - high performance routing 614
 - index 521
 - LU-LU session 556
- network
 - deactivation and session termination 492
 - description 447
 - error detection and recovery and SSCP management services 512
 - generic BIND AMRUs 451
 - index 447
 - resource activation 453
 - session establishment 466
- FNDADJCP VTAMMAP 222
- FNDANDCB VTAMMAP 223
- FNDNDCB VTAMMAP 224
- FNDDECB VTAMMAP 225
- FNDENDEL VTAMMAP 226
- FNDLDCB VTAMMAP 227
- FNDNDREC VTAMMAP 228
- FNDNDWGT VTAMMAP 230
- FNDNODE VTAMMAP 231
- FNDREREC VTAMMAP 232
- FNDSCCB VTAMMAP 233
- FNDSTICB VTAMMAP 234
- FNDTGREC VTAMMAP 238
- formatting and printing dump output 188
- formatting trace output, using ACF/TAP 323
- full-screen
 - description of mode (TSO/VTAM) 113
 - incorrect processing (TSO/VTAM) 105
- function error (TSO/VTAM) 113

G

- generalized PIU trace (GPT)
 - description 353
 - operation 354
 - when to use 354
- generalized trace facility (GTF) 322
- GPT (generalized PIU trace)
 - description 353
 - operation 354
 - when to use 354
- Greenwich mean time (coordinated universal time) 275
- GTF (generalized trace facility) 322
- GUNBIND pend, cross domain log off 15

H

- hard-copy log 51
- hardware
 - error recording 403
 - error records, LOGREC 51
 - problem documentation 49
- High-Performance Routing flows 614
- hold option, effect on performance 119
- host IRN (ISTIRN), I/O trace for 341
- host physical unit (ISTPUS)
 - buffer contents trace for 331
 - I/O trace for 341

- HOST VTAMMAP 240, 241
- HPR route test 150
- hung
 - session 64
 - terminal (TSO/VTAM)
 - symptoms 105

I

- I/O pending 69
- I/O trace
 - description 340
 - operation 341
 - when to use 340
- I/O traces 322
- IBM Software Support Center, contacting xviii
- IBM Support Center 103
- IBM Support Center, documentation for
 - general information 55
- IBMLINK 4
- IBMTTEST command (logical unit connection test) 403
- IKJ608I message 105
- IKT028I message 105
- IKT029I message 105
- IMR (intensive mode recording) 179
- incorrect
 - data
 - length (TSO/VTAM) 105
 - translation (TSO/VTAM) 116
 - line prompting (TSO/VTAM) 105
 - output
 - diagnosis procedure 90
 - symptoms 4, 42
 - TSO/VTAM 112
 - processing for a mode (TSO/VTAM) 105
 - screen management (TSO/VTAM) 113
 - screen size
 - diagnosis procedure (TSO/VTAM) 117
 - documentation requirements (TSO/VTAM) 117
 - non-SNA 3270 terminal 117
- increase in
 - detected waits 106
 - swap-outs 106
- indirect address word structure 437
- Information APARs xxi
- input is concatenated 115
- input mode processing is incorrect (TSO/VTAM) 105
- input/output problem determination (IOPD) facility 177
- installing dump analysis and trace analysis tools 641
- intensive mode recording (IMR) 179
- intermediate routing node (IRN), I/O trace of 341
- Internet, finding z/OS information online xxiii
- interpreting SVC trace entries 105
- IOPD (input/output problem determination) facility 177
- IRN, trace of (host) 341
- IST1774I message 30
- IST1775I message 30
- IST1831I 98
- IST1832I 98
- IST1833I 98
- IST562I 98
- IST564I 98
- ISTEXCAA, Abend 9
- ISTEXCCS, Abend 9
- ISTEXCSD, Abend 9
- ISTEXCVR, Abend 9
- ISTPATCH module 405

- ISTPUS (host physical unit)
 - buffer contents trace for 331
 - I/O trace for 341
- ISTPUS in traces
 - buffer trace 331
 - I/O trace 341
- ISTTSCPF 8
- ISTVABND VTAMMAP 242
- ISTVDUMP VTAMMAP 244
- ISTVMAP VTAMMAP 247
- ISTVSAVE VTAMMAP 249
- ISTVSLIP VTAMMAP 251

J

- JOBN trace field 325

K

- keyboard 655
- keyboard, character functions improperly 115

L

- library, VTAM definition 53
- license, patent, and copyright information 659
- line delete key functions improperly 115
- line mode, description (TSO/VTAM) 113
- line trace
 - description 354
 - operation 355
 - record format
 - communication scanner type 2 355
 - communication scanner type 3 357
 - when to use 355
- link level 2 (LL2) test 49, 179
- link pack area (LPA) map 51
- link-edit (XREF) map 50
- LL2 (link level 2) test 49, 179
- local SNA terminals, pacing values for (TSO/VTAM) 120
- location of BIND 118
- locked queue anchor block (LQAB) 69
- locks, description of 74
- log of console 52
- logical unit (LU) hangs, PNFYx state 9
- logical unit connection test (IBMTTEST) 403
- logmode
 - default name in USS command 117
 - table entry in DLOGMOD 107
- logon problems
 - CICS 17
 - TSO/VTAM
 - ABEND0AB 108
 - cross-domain network 108
 - diagnosis procedure 106
 - documentation requirements 106, 119
 - fails for all terminals 108
 - symptoms 105
- LOGREC error recording data set 51
- loop problems
 - diagnosis procedure 81
 - symptoms 4, 42
- LPA (link pack area) map 51
- LPDA-2, network flows 514, 515, 516
- LQAB control blocks 69
- LU (logical unit) hangs, PNFYx state 9

- LU-LU session flows 556

M

- mainframe
 - education xxi
- maintenance and operator subsystem (MOSS) dump 188
- MAXDATA value, how to specify 114
- MDR (miscellaneous data record) 403
- messages
 - issued for 3745 bus switching 402
 - IST1278I
 - GUNBIND PENDING 15
 - IOPD facility, issued from 178
 - NMVT PENDING 16
 - IST1774I 30
 - IST1775I 30
 - IST1832I 98
 - IST1883I 98
 - IST259I 10
 - IST400I 108
 - IST530I 16
 - IST562I 98
 - IST564I 98
 - IST718I 18
 - IST719I 18
 - IST804I 108
 - IST805I 108
 - module identification, modify 178
 - problem
 - procedure 87
 - symptoms 4, 42
 - TSO/VTAM logon problems 105
- miscellaneous data record (MDR) 403
- mode switching errors (TSO/VTAM) 105
- MODEENT macro, description 107
- modified data tags 115
- MODIFY command
 - clear EE connection network unreachable partner information 180
 - dump 186
 - IOPD
 - enabling the IOPD facility 177
 - IOMSGLIM start option, effect on IOPD facility 178
 - message module identification 178
 - MSGMOD 178
 - NCP intensive mode recording 179
 - SDLC link level 2 test 179
 - trace 308
 - tuning statistics 181
 - VTAMOPTS 181
- MODIFY CSDUMP command 177
- module name, finding in a dump 59
- module trace 312
- MOSS dump 188
- MVS
 - dumps
 - abend 183
 - formatting and printing 188
 - stand-alone 184
 - SVC 184
 - VTAM control blocks formatted 189
 - performance group specification 120
 - trace fields 325

N

NCP (Network Control Program)

- dumps
 - dynamic 185
 - independent dump utility 187
 - static 185
 - VTAM dump facility 186
- error recording 404
- intensive mode recording 179
- storage, displaying 155
- traces
 - generalized PIU (GPT) 353
 - line 354
 - network controller line 358
 - scanner interface (SIT) 359
 - transmission group (TG) 359

NetView program

- file 51
- for hardware failure 49
- hard-copy log 51

NetView session trace 313

Network Control Program (NCP)

- dumps
 - dynamic 185
 - independent dump utility 187
 - static 185
 - VTAM dump facility 186
- error recording 404
- intensive mode recording 179
- storage, displaying 155
- traces
 - generalized PIU (GPT) 353
 - line 354
 - network controller line 358
 - scanner interface (SIT) 359
 - transmission group (TG) 359

network controller line trace

- operation 358
- output 359
- when to use 358

network status, displaying 156

NMVT alerts, error recording 401

non-SNA 3270 terminal, incorrect screen size 117

NVPACE operand (TSO/VTAM) 120

O

OBR (outboard recorder) record 403

- operation checks 105
- operation codes, CPCB 623

operator commands

- DISPLAY 123
- MODIFY 123

option, TPUT, location of 115

outboard recorder (OBR) record 403

output

- problems, TSO/VTAM 112
- wait state 120

P

PABs (process anchor blocks) 66

PABSCAN VTAMMAP 257

pacing, values for local SNA terminals (TSO/VTAM) 120

panel interface, using the 193

PARTNRLU VTAMMAP 261

patch area

- TSO/VTAM 405
- VTAM 404

path problems, example solution 157

path tables, display of 155

PDDDB, problem determination database 4

pending state, display of resources in 156

performance group

- how to specify 120
- location of 120

performance problem

- diagnosis procedure 94, 119
- symptoms 4, 42
- TSO/VTAM 106

PIU too long 114

PNFYx state 9

prerequisite information xxi

primary screen size, PSERVIC coding for 118

printing output

- dumps
 - using ABDUMP (MVS) 189
 - using SADMP (MVS) 190
- traces
 - 3705, 3720, 3725, or 3745 SMS line trace 355
 - using ACF/TAP 323

problem determination commands

DISPLAY

- buffer pool use 127
- device workload information 175
- HPR route test 150
- NCP storage 155
- network status 156
- path tables 155
- pending state 156
- route status 156
- route test 163
- topology database update information 167
- traces 175
- VTAM storage 175

MODIFY

- intensive mode recording 179
- IOPD 177
- link level 2 test 179
- message module identification 178
- tuning statistics 181

problem determination database (PDDDB) 4

problem documentation 48

problem number 55

problem types

- abend (abnormal end) 57
- documentation 100
- failing module 100
- incorrect output 90
- loop 81
- message 87
- path, example solution 157
- performance 94
- storage 97
- VTAM vs. non-VTAM 3
- wait, DSRLIST 16

problems and symptoms, common

APPN networks

- descriptions 26
- index of problems 22

documentation

- submitting 54

- problems and symptoms, common (*continued*)
 - documenting
 - recommended 48
 - HPR networks
 - descriptions 38
 - index of problems 38
 - subarea network
 - description 8
 - VIT analysis tool
 - documenting 44
 - isolating 44
 - VTAM dump analysis tool
 - documenting 48
 - isolating 47
 - VTAM, index of problem types 42
- process anchor blocks (PABs) 66
- process scheduling table (PST) 65
- processing is hung (TSO/VTAM) 105
- program
 - temporary fix (PTF) 52
 - update tape (PUT) 52
- program temporary fix (PTF) eyecatcher 52
- PROGxxx message, TSO/VTAM 105
- PST (process scheduling table) 65
- PTF (program temporary fix) eyecatcher 52

Q

QDIOSYNC trace 341

R

- RDTCHECK VTAMMAP 262
- RDTFULL VTAMMAP 263
- RDTHIER VTAMMAP 264
- RDTSUM VTAMMAP 265
- RECMS (record management statistics) 11
- record management statistics (RECMS) 11
- repeated error record entries, LOGREC 85
- reporting problems procedure 103
- request parameter header (RPH)
 - finding 73
 - waiting 73
- reshow processing, description 113, 115
- resource state trace 345
- response time slow (TSO/VTAM) 106
- REXX exec for VIT analysis tool 367
- RFC (request for comments)
 - accessing online xxiii
- route
 - status, display 156
 - test, display 163
- ROUTES VTAMMAP 266
- RPH (request parameter header)
 - finding 73
 - waiting 73
- RPL request type, location in dump of SDWA 109
- RPLFDB2, location in dump of SDWA 109
- RPLFDBK2, location in dump of SDWA 109
- RPLRTNCD, location in dump of SDWA 109
- RTPINFO VTAMMAP 267

S

sample procedures 193

- save area
 - module linkage conventions, APPN 407
 - module linkage conventions, subarea 405
- SAW data, filtering 16
- scanner interface trace (SIT)
 - description 359
 - operation 359
 - when to use 359
- screen management problems, TSO/VTAM
 - diagnosis procedure 112
 - documentation 112
 - symptoms 105
 - types 112
- screen size
 - changing, in non-full-screen processing 118
 - default value 118
 - problems, TSO/VTAM 115, 117
 - PSERVIC coding 118
 - TSO/VTAM problems 106
- SCRSIZE operand 118
- SDLC link level 2 test, modify 179
- SDWA (system diagnostic work area) 52
- sense codes 41
 - 0801 20
 - 0821 20
 - 0835 20
 - 800A 21, 41, 42
- sense data in transmission group trace, MVS 361
- service aids, diagnosis tools
 - alert messages 401
 - APPN 407
 - hardware error recording 403
 - IPCS 188
 - logical unit connection test (IBMTTEST) 403
 - messages for 3745 bus switching 402
 - NCP error recording 404
 - NCP intensive mode recording 179
 - patch area 404
 - recording NMVT alerts in LOGREC 401
 - save area module linkage conventions, subarea 405
- SES VTAMMAP 271
- session
 - awareness (SAW) data 51
 - ending unexpectedly 11
 - parameters, defining (TSO/VTAM) 107
 - trace data 51
- session awareness data, filtering 16
- session management exit (SME) buffer trace 346
- shortcut keys 655
- SIBCHECK VTAMMAP 275
- SIT (scanner interface trace)
 - description 359
 - operation 359
 - when to use 359
- slow response time
 - general procedure 94
 - TSO/VTAM 106
- SMS (buffer use) trace
 - description 348
 - format, MVS 349
 - operation 349
 - when to use 348
- SNA protocol specifications 653
- SNA, terminals, local, pacing values for (TSO/VTAM) 120
- SNAPREQ start option 96
- softcopy information xxi
- solving problems, where to begin 3

- SPANC VTAMMAP 277
- SRTFIND VTAMMAP 281
- stand-alone dump, MVS 184
- START command, using to activate VTAM traces 308
- states of resources, tracing 345
- static NCP dump 185
- status, display of 156
- storage and control block ID codes 637
- storage problem 18, 97
- STORAGE VTAMMAP 282
- STSIZE macro, changing screen size with 118
- subplex
 - MNPS session recovery error 37
 - VTAM locks 74
- Support Center 103
- SVC dump 184
- swap count, incorrect incrementing 120
- swap-outs increase 106
- symptom string
 - description 52
 - MVS 58
 - structure 102
- symptoms
 - abend 4, 57
 - abend, in TSO/VTAM 109
 - documentation problem 100
 - incorrect output 90
 - logon problem, in TSO/VTAM 106
 - loop problem 81
 - message problem 87
 - performance problem 94
 - performance problem, in TSO/VTAM 106
 - screen management problems 105
 - screen size problems 106
 - storage problem 97
- syntax diagram, how to read xix
- syntax diagrams, reading 191
- SYS1.LOGREC error recording data set 51
- system diagnostic work area (SDWA) 52

T

- TAP (trace analysis program) 323
- tasks
 - accessing VTAM formatted dump
 - steps for accessing VTAM formatted dump 193
 - obtaining PUTDOC
 - steps for 53
 - obtaining TRSMAN
 - steps for 54
 - using the batch option
 - steps for using the batch option 195
 - using the IPCS command line
 - steps for using the IPCS command line 194
- TCP/IP
 - online information xxiii
- Technotes xxi
- terminal
 - cannot log on 106
 - definition statement 107
 - device problem 85, 92
 - incorrect output problem 117
 - name, location in dump of SDWA 109
 - pacing values 120
 - user echo test 403
- termination, TSO/VTAM
 - abend 109

- termination, TSO/VTAM (*continued*)
 - during logon 106
- TNSTAT (tuning statistics), modify 181
- topology database update statistics, displaying 167
- TOPOLOGY VTAMMAP 283
- TPUT option
 - editing 116
 - location 115
- trace analysis tools, installing dump analysis and 641
- traces
 - activating 308
 - analysis program (ACF/TAP) 323
 - buffer contents 328
 - buffer use (SMS) 348
 - CCW 92, 310
 - component execution sequence 405
 - fields, MVS-only 325
 - generalized PIU (GPT) 353
 - I/O 340
 - line 354
 - network controller line 358
 - resource state trace 345
 - scanner interface (SIT) 359
 - session management exit (SME) buffer 346
 - SMS (buffer use) 348
 - TGET/TPUT, for TSO/VTAM 350
 - tool, VIT analysis 365
 - transmission group (TG) 359
 - using 307
- trademark information 667
- translation
 - incorrect (TSO/VTAM) 116
 - tables (TSO/VTAM) 117
- transmission group (TG) trace
 - description 359
 - operation 360
 - output for MVS 361
 - when to use 360
- TRSTRACE VTAMMAP 285
- TSO EDIT problems 105
- TSO TERMINAL command, changing screen size with 118
- TSO/VTAM problems
 - ABEND0AB 109
 - ABEND0AC 110
 - ABEND0AD 110
 - ABEND15D 110
 - data misplaced on screen 115
 - data translated incorrectly 116
 - exception responses 114
 - extra data 114
 - first logon from a particular device fails 107
 - first logon using USS commands fails 106
 - function error 113
 - incorrect data translation 116
 - incorrect screen size 117
 - initial TSO/VTAM problem analysis 105
 - logon fails 10, 106
 - misplaced data on screen 115
 - missing data 114
 - mode error (incorrect screen management) 113
 - parameter initialization 110
 - performance 119
 - response time is slow 106
 - screen is wrong size for mode 118
 - screen management 112
 - TGET, partial input for 19
 - translation of data 116

TSO/VTAM TGET/TPUT trace
 description 350
 operation 351
 output 352
 when to use 351
 tuning statistics (TNSTAT), modify 181

U

user cannot log on TSO/VTAM 106
 user edit exits, TSO/VTAM
 when to use 116
 where to find description of 115
 where to find list of 117
 user exit, Abend in 9
 USS commands used in logon 106
 USS message 10 105
 USS message 7 105
 utility
 IEBCGENER 55

V

VERBEXIT VTAMMAP, IPCS subcommand
 ALL 196
 APPLCONV 197
 APPLMODAL 200
 APPLMODE 198
 APPNBASE 202
 FINDDSIB 216
 FINDSIB 218
 FNDADJCP 222
 FNDANDCB 223
 FNDCOS 224
 FNDDECB 225
 FNDENDEL 226
 FNDLCB 227
 FNDNDREC 228
 FNDNDWGT 230
 FNDNODE 231
 FNDREREC 232
 FNDSCCB 233
 FNDSITCB 234
 FNDTGWGT 238
 HOST 240, 241
 PABSCAN 257
 PARTNRL 261
 RDTCHECK 262
 RDTFULL 263
 RDTHIER 264
 RDTSUM 265
 ROUTES 266
 SES 271
 SIBCHECK 275
 SPANC 277
 SRTFIND 281
 STORAGE 282
 VITAL 286
 VTAM 287
 VTBASIC 289
 VTBUF 290
 VTCVTPAB 291
 VTFNDMOD 294
 VTMODS 295
 VTNODE 297
 VTREADYQ 298

VERBEXIT VTAMMAP, IPCS subcommand (*continued*)
 VTRPH 299
 VTVIT 300
 VTVRBLK 302
 VTWRE 304
 VIT analysis tool
 data set, required to run 366
 description 365
 documenting 47
 I/O options 397
 interactive VIT analysis 367
 isolating 44
 job parameters 368
 job submission 369
 output, checking 369
 parameter dataset, creating 399
 problems 44
 REXX exec 367
 setting up and running 365
 timing options 394
 VIT recording 366
 VITAL VTAMMAP 286
 VTAM
 buffer pool use 348
 channel programs 413
 commands for problem determination 124
 control block ID codes 637
 control blocks formatted in a dump, MVS 189
 definition library 53
 dump analysis, using IPCS VTAMMAPs 191, 192
 dump facility 186
 execution sequences 405
 locks 79
 service aids 183, 307, 365
 trace fields, MVS-only fields 325
 traces
 activating 308
 buffer contents 328
 buffer use (SMS) 348
 I/O 340
 printing with ACF/TAP 323
 resource state 345
 session management (SME) exit buffer trace 346
 SME buffer 346
 SMS (buffer use) 348
 TGET/TPUT, for TSO/VTAM 350
 wait state indications 4, 42
 VTAM dump analysis, diagnosing problems in 47
 VTAM internal trace (VIT)
 extracting information
 address 388
 boolean expression freeform 389
 description 385
 offset, address 388
 offset, string, character 388
 offset, string, hex 388
 operands, boolean freeform 389
 options, entries 387
 sample output 392
 string, character 388
 string, hex 388
 syntax 391
 template 386
 missing trace records 19
 RU counting
 network address 378
 output, sample 382

- VTAM internal trace (VIT) *(continued)*
 - RU counting *(continued)*
 - parameters, RU 379
 - request/response units 377
 - RU 378
 - syntax 378
 - storage, analyzing
 - ASID 371
 - buffer pools 371
 - description 369
 - matching 370
 - sample output 372
 - storage lengths 370
 - syntax 371
 - unmatched allocate 370
- VTAM locks, description of 74
- VTAM VTAMMAP 287
- VTAM, online information xxiii
- VTAMMAP control card 188
- VTBASIC VTAMMAP 289
- VTBUF VTAMMAP 290
- VTCVTPAB VTAMMAP 291
- VTFNDMOD VTAMMAP 294
- VTMODS VTAMMAP 295
- VTNODE VTAMMAP 297
- VTREADYQ VTAMMAP 298
- VTRPH VTAMMAP 299
- VTVIT VTAMMAP 300
- VTVRBLK VTAMMAP 302
- VTWRE VTAMMAP 304

W

- wait
 - application program 64
 - conditions
 - process waits for a buffer 73
 - process waits for a resource 73
 - option 119
 - session 64
- waiting
 - request element (WRE) 69, 71
 - RPH 73
- work areas (to trace execution sequences) 405
- wrapping of data on screen 115
- WRE (waiting request element) 69, 71
- write control character in buffer contents trace output 112

X

- XREF (link-edit) map 50

Z

- z/OS Basic Skills Information Center xxi
- z/OS, documentation library listing 669

Communicating your comments to IBM

If you especially like or dislike anything about this document, use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this document. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

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

Send your comments to us in any of the following ways:

- To send comments by FAX, use this number: 1+919-254-1258
- To send comments electronically, use this address:
 - comsvrcf@us.ibm.com
- To send comments by post, use this address:

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

Make sure to include the following information in your note:

- Title and publication number of this document
- Page number or topic to which your comment applies



Product Number: 5650-ZOS

Printed in USA

GC27-3667-00

