

z/OS Communications Server
3.2

SNA Network Implementation Guide



Note:

Before using this information and the product it supports, be sure to read the general information under [“Notices” on page 629](#).

This edition applies to 3.1 of z/OS® (5655-ZOS), and to subsequent releases and modifications until otherwise indicated in new editions.

Last updated: 2025-09-20

© **Copyright International Business Machines Corporation 2000, 2025.**

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

Contents

Figures.....	xv
Tables.....	xxiii
About this document.....	xxvii
Who should read this document.....	xxvii
How this document is organized.....	xxvii
How to use this document.....	xxviii
How to provide feedback to IBM.....	xxviii
Conventions and terminology that are used in this information.....	xxviii
How to read a syntax diagram.....	xxx
Prerequisite and related information.....	xxxii
Summary of changes for SNA Network Implementation Guide.....	xxxvii
Summary of changes for z/OS 3.2.....	xxxvii
Changes made in z/OS Communications Server 3.1.....	xxxvii
Chapter 1. Post-installation considerations for z/OS Communications Server.....	1
Defining Communications Server SNA to z/OS.....	1
Using automatic restart manager.....	1
Starting z/OS Communications Server SNA.....	2
Chapter 2. VTAM networking concepts.....	3
VTAM functions.....	3
VTAM nodes.....	3
Nodes with APPN function only.....	3
Nodes with subarea function only.....	6
Nodes with both subarea and APPN function.....	7
Network accessible units.....	9
Physical unit.....	9
Logical unit.....	9
Network control sessions.....	10
SSCP-SSCP sessions.....	10
CP-CP sessions.....	11
User sessions.....	11
How VTAM locates resources.....	12
Locating resources in a subarea network.....	12
Locating resources in an APPN network.....	13
Route selection.....	13
Routing in a subarea network.....	13
Routing in an APPN network.....	14
Addressing.....	16
Subarea addressing.....	16
APPN addressing.....	16
Controlling network data flow using pacing.....	16
Pacing concepts.....	17
Pacing window.....	17
Chapter 3. Implementing a VTAM network.....	19

Using start options and configuration lists.....	20
Start options.....	21
Configuration lists.....	28
Identifying resources to VTAM.....	29
Coding concepts.....	29
Sift-down effect.....	30
Using MVS system symbols.....	30
Verifying a VTAM network.....	36
Verifying a multiple-domain subarea network.....	37
Verifying a multiple-network environment.....	37
Verifying a VTAM APPN network.....	37
Chapter 4. Connecting an APPN node to VTAM.....	39
Connections through boundary function-based transmission groups.....	39
Multiple connections with parallel transmission groups.....	40
Channel connections between APPN nodes.....	40
Multipath channel connections.....	40
Composite network node channel connections.....	46
Leased connections between APPN nodes.....	47
IBM 3172 Nways Interconnect Controller connections between APPN nodes.....	48
Examples.....	49
Using a connection network.....	50
IBM Open Systems Adapter connections between APPN nodes.....	54
APPN multiple network connectivity.....	75
Peripheral subnetwork boundaries.....	77
Extended subnetwork boundaries.....	77
APPN multiple network connectivity support.....	77
Virtual-route-based transmission groups.....	80
Defining a VR-based TG.....	82
VR-based TG recommendations.....	83
Selecting the network node server for end nodes.....	85
Creating a network node server list.....	85
Activating, replacing, and displaying a network node server list.....	86
Using the NNSPREF start option.....	86
Chapter 5. Connecting a subarea node to VTAM.....	87
Connecting two VTAMs using channels.....	87
Channel-to-channel adapter connection.....	87
Multipath channel connections.....	89
Chapter 6. Using Enterprise Extender (EE).....	91
Overview.....	92
Benefits of Enterprise Extender.....	92
Availability of Enterprise Extender.....	92
Hardware requirements.....	92
EE reliability and strategy.....	93
Using EE and extended border node (EBN) as a replacement for SNI.....	94
EE implementation considerations.....	94
Designing the EE network.....	94
Distinctions between an EE network and an EE connection network.....	95
Characteristics of EE connections.....	95
Static VIPA considerations.....	96
IP multipath considerations.....	97
Class of Service preservation dependencies.....	97
Comparing host name and IP address definitions.....	98
Configuring the EE network.....	99
Preservation of SNA transmission priority.....	101

Network address translation (NAT) considerations	101
Steps for configuring and activating an EE network.....	101
Configuring the EE connection network.....	116
Connection network concepts.....	121
EE connection network rules.....	122
Contrasting local and global networks.....	122
Benefits of defining multiple Enterprise Extender virtual routing nodes.....	122
Defining an EE connection network in the EE XCA major node	125
EE security considerations.....	127
SNA session level encryption (SLE)	127
IP security (IPSec).....	127
Using EE with network address translation (NAT).....	128
IP filtering.....	129
IDS for Enterprise Extender.....	129
OEM security products - EE proxy solutions.....	130
Tuning the EE network.....	130
Tuning Enterprise Extender-specific buffer pools	130
Timers.....	131
HPR ALIVE timer optimization for Enterprise Extender.....	134
Enterprise Extender LDLC keep-alive reduction.....	134
Enterprise Extender improved packet loss tolerance.....	135
Disconnect and inactivity summary.....	135
Customizing IP type of service.....	136
Advanced coding considerations for EE.....	137
EE connection network reachability awareness	137
TCP/IP MTU size for EE	144
Running EE in constrained or virtualized environments.....	145
RTP transmission stall operator awareness and recovery support.....	146
Load balancing.....	147
Transmission group profiles.....	147
Dynamic reconfiguration.....	148
Dial usability - DWACT, DWINOP, KEEPACT, REDIAL, and REDDELAY	148
Customization for EE connection network PUs.....	149
Cross-subnet routing with global VRNs.....	150
Troubleshooting EE problems.....	152
Chapter 7. OSA-Express.....	159
OSA-Express overview.....	159
Defining an OSA-Express device to z/OS Communications Server using QDIO.....	159
OSA routing.....	161
OSA-Express virtual MAC (VMAC) routing	161
Primary and secondary routing.....	163
Outbound priorities.....	165
MTU.....	165
Chapter 8. Defining resources dynamically.....	167
Defining switched resources dynamically.....	167
Dynamic PU definition (DYNPU operand).....	167
Dynamic switched definitions.....	167
Dynamic configuration of channel-attached devices.....	170
Installation and preparation.....	171
Defining your configuration.....	172
Building resource definitions.....	172
Using the default naming convention.....	173
Customizing the command lists.....	173
Dynamic reconfiguration and change of operands.....	174
Dynamic reconfiguration and dynamic change requirements.....	175

Using the VARY ACT,UPDATE technique.....	176
Using the VARY DRDS technique.....	178
Using the MODIFY DR technique.....	180
Dynamic reconfiguration of independent LUs.....	180
Chapter 9. Defining peripheral nodes.....	183
Defining type 2.1 peripheral nodes.....	183
Nonnative network type 2.1 connections.....	184
Attaching peripheral nodes to VTAM.....	184
Local non-SNA connection.....	185
Local SNA connection.....	185
Loop-adapter-attached connection.....	187
Chapter 10. Defining LUs.....	189
Independent LUs.....	189
Characteristics of independent LUs.....	189
Defining independent LUs.....	190
Multiple connections between a type 2.1 node and a subarea node.....	194
Dynamic selection of session connections.....	195
Authorized transmission priority for LEN connections.....	196
Restrictions on using independent LUs.....	196
Dependent logical units.....	196
Chapter 11. Establishing and controlling SNA sessions.....	199
Multicultural support.....	199
Multicultural support for user USS messages and commands.....	200
Defining USS tables for use with the LANGTAB operand	200
Defining USS messages to the MVS message service.....	200
Multicultural support for the language passed to application programs.....	201
Model name table.....	201
Associated LU table.....	202
Selecting session parameters for the logon mode table.....	203
Automatic logons.....	209
Coding for automatic logon.....	210
Operator commands for automatic logon.....	212
Reallocation of autologon sessions.....	212
Session management exits.....	215
Session authorization.....	216
Session accounting.....	217
Session-level pacing.....	217
Fixed session-level pacing.....	218
Adaptive session-level pacing.....	218
Setting initial pacing values.....	219
Sample configurations.....	224
Logon and logoff requests from dependent logical units.....	232
Unformatted logon requests using mixed-case passwords.....	233
Chapter 12. Network routing.....	235
Network routing and resource location for APPN nodes.....	235
Types of searches.....	236
Minimizing broadcast searches.....	238
Network search overhead caused by duplicate resource definitions.....	241
Avoiding congestion.....	241
Checkpointing of the TRS database and the directory database.....	242
APPN Class of Service.....	242
APPN network routing through a composite network node (CNN).....	250
Using the SAMAP table.....	253

Network routing for subarea nodes.....	256
Physical paths.....	256
Logical paths.....	258
How session traffic is assigned to a specific route.....	260
How to plan routes in your network.....	264
How VTAM handles network and subarea addressing.....	267
Virtual route pacing.....	270
Parallel sessions using parallel transmission groups.....	273
Dynamic path update.....	274
Chapter 13. Application programs.....	277
Naming an application program.....	277
Model application program definitions.....	278
Overview.....	278
Coding guidelines.....	279
How VTAM finds the best match.....	281
Example of using model application program definitions.....	282
Resource state requirements.....	283
Authorizing application facilities.....	283
Passing and validating logon requests.....	283
Overriding dial number digits for dial connections.....	283
Acquiring LU sessions.....	284
Enabling parallel sessions.....	284
Authorizing privileged paths.....	284
Data compression.....	284
Types of compression.....	285
Implementing data compression.....	285
Compression level negotiation.....	286
Summary of data compression.....	288
Security features.....	288
Cryptography facility.....	289
Message authentication.....	297
SLU authentication.....	298
VTAM application security.....	298
Confidential data.....	299
3270 Intrusion Detection Services.....	299
Logon mode parameters.....	318
Using user variables (USERVAR).....	319
Application workload balancing with USERVAR.....	319
Classes of USERVARs.....	319
Types of USERVARs.....	320
Processing USERVARs.....	321
USERVAR propagation and routing.....	322
Defining your network with USERVARs.....	323
Dynamic USERVAR update session failure.....	323
Generic resources function.....	324
High availability using extended recovery facility.....	325
Security features in an XRF environment.....	326
LU 6.2 in an XRF Environment.....	326
Persistent LU-LU sessions.....	327
Single node persistent sessions.....	328
VTAM common network services.....	329
Cross-memory application programming interface (API).....	329
Allocating private storage.....	329
Communicating with start-stop devices.....	329
LU 6.2 application programs.....	330
Enabling LU 6.2 support.....	330

LU 6.2 sessions.....	330
LU 6.2 session limits.....	332
Managing LU 6.2 sessions with operator commands.....	333
LU 6.2 security.....	334
LU 6.2 sync point services.....	335
Selective termination of idle LU 6.2 sessions.....	335
Selective termination of idle network management sessions.....	336
APING support.....	337
High-performance data transfer (HPDT).....	339
Communications storage manager (CSM).....	340

Chapter 14. Functions provided by VTAM in a sysplex..... 341

Setting up the sysplex environment for VTAM and TCP/IP functions.....	341
Sysplex subplexing.....	341
Considerations.....	343
Coupling facility structure attributes.....	344
Determining the size of the coupling facility structure.....	346
Sample CFRM coding.....	346
Connecting to and allocating storage for coupling facility structures.....	348
Structure rebuild.....	348
Coupling facility duplexing.....	349
Coupling facility storage shortages.....	349
Dynamic altering of structures.....	349
Dynamic definition of VTAM-to-VTAM connections.....	350
Generic resources.....	352
Generic resources requirements.....	352
Generic resource mapping.....	353
Partner LU mapping.....	354
Initiating sessions with generic resource members.....	357
Implementation considerations.....	358
Coupling facility failures for generic resource configuration.....	359
Removing a generic resource.....	360
Routine maintenance for VTAM nodes.....	365
Multinode persistent sessions.....	366
Multinode persistent session configuration requirements.....	368
Using multiple coupling facility structures for multinode persistent sessions.....	369
Establishing multinode persistent sessions.....	370
Coupling facility failures for multinode persistent session configuration.....	373
Failure recovery processing.....	374
MNPS planned and forced takeover processing.....	376
What to do if recovery does not occur or complete.....	379
TSO generic resources.....	380
Sysplex-wide security associations.....	380
Coupling facility failures for sysplex-wide security associations.....	381
Failure of a TCP/IP stack.....	381
Failure of a VTAM node.....	381
Rebuild of the sysplex-wide security associations structure (EZBDVIPA).....	381
Disconnect from the EZBDVIPA structure.....	381
Modifying the number of lists.....	382
Sysplexports.....	383
Coupling facility failures for Sysplexports.....	383
Failure of a TCP/IP stack.....	383
Failure of a VTAM node.....	383
Rebuild of the Sysplexports structure (EZBEPOR).....	383
Disconnect from the EZBEPOR structure.....	384

Chapter 15. Implementing an APPN network.....385

Coding considerations for APPN resources.....	385
Maximum APPN Locate size considerations.....	386
Enabling control sessions.....	387
CP-CP sessions between two VTAM nodes.....	387
Defining adjacent APPN nodes.....	387
Defining the logon mode for CP-CP sessions.....	388
High-Performance Routing (HPR).....	388
What is High-Performance Routing?.....	388
What is Rapid Transport Protocol?.....	389
What is automatic network routing?.....	390
How does HPR switch paths?.....	390
HPR implementation overview.....	391
Setting session paths using HPRNCPBF.....	399
Chapter 16. Implementing a combined APPN and subarea network.....	401
Start options defining a combined subarea and APPN environment.....	404
SORDER start option.....	404
SSEARCH start option.....	404
CDRSCTI start option.....	404
IOPURGE start option.....	405
Dependent LUs.....	405
Dependent LU server.....	405
Dependent LU server support across subnetwork boundaries.....	408
Defining CDRSCs.....	410
SSCP takeover.....	410
Establishing and controlling sessions.....	412
Controlling searches.....	413
Using SORDER to control network search order.....	413
Using SSEARCH to limit subarea network searches.....	414
Using the CDRSCTI start option to reduce broadcast searches of APPN.....	414
Using the DISJOINT operand to define disjoint subarea networks.....	414
APPN and subarea Class of Service resolution.....	416
Resolving logon mode names to subarea and APPN Classes of Service.....	417
Defining APPN and subarea Classes of Service in logon mode tables.....	419
Defining APPNTOSA and SATOAPPN class of service mapping tables.....	420
Adding and moving nodes.....	420
Chapter 17. Implementing a subarea network.....	421
Start options defining other domains.....	422
Defining the location of cross-domain resource managers by coding adjacent SSCP tables.....	422
Defining the location of cross-domain resource managers dynamically.....	423
Specifying timeout values for locating cross-domain resources.....	423
Identifying VTAMs in other domains (CDRMs).....	423
Identifying resources in other domains.....	425
Dynamic definition of cross-domain resources.....	425
Static definition of cross-domain resources.....	427
Model definition of cross-domain resources.....	428
Adjacent SSCPs.....	431
CDRM owner verification for cross-domain resources.....	437
Changing ownership of cross-domain resources.....	438
Shadow resources.....	438
Chapter 18. Connecting multiple subarea networks.....	439
Defining a multiple-network environment.....	439
SNI configurations.....	441
Start options defining other networks.....	442
Start options for gateway VTAMs.....	442

GWSSCP start option for nongateway VTAMs.....	442
Configuration lists for gateway VTAMs.....	442
Connecting networks.....	443
Defining a gateway VTAM.....	444
Defining cross-domain resource managers.....	444
Defining cross-domain resources.....	447
Session initiation request.....	448
Name assumption.....	449
Predefined cross-domain resources without network specification.....	450
Predefined cross-domain resources with network specification.....	451
Dynamic cross-domain resources.....	452
Defining adjacent SSCPs.....	454
Types of adjacent SSCP tables.....	454
Deciding whether to code adjacent SSCP tables.....	454
Sample of adjacent SSCP tables for a multiple-network environment.....	454
Request routing.....	457
Dynamically defined CDRSCs and adjacent SSCP tables.....	458
Alias name translation and adjacent SSCP tables.....	458
Cross-network routing.....	458
Network address structures.....	459
Network naming conventions.....	459
Controlling paths for interconnected networks.....	460
Handling class of service tables.....	461
Address translation.....	463
Resource name translation.....	465
Alias selection function of the session management exit routine.....	466
NetView alias name translation facility.....	466
Establishing and controlling SNA sessions.....	472
Nonnative network type 2.1 connections.....	472
Automatic logon.....	473
Operating VTAM.....	474
Using the NetView program for network management.....	474
Application programs.....	474
Defining the NetView program.....	474

Chapter 19. Operating VTAM..... 477

Starting the domain.....	477
Configuration restart.....	477
Information recorded by configuration restart.....	481
Activating resources.....	481
Order of activation.....	482
Resources automatically activated by VTAM.....	483
Activating application programs.....	486
Monitoring the domain.....	486
Using the DISPLAY command.....	486
Monitoring I/O problems.....	486
Suppressing messages.....	486
Message flooding prevention.....	487
Other methods of controlling messages.....	487
Displaying and testing routes.....	488
Defining operator messages and commands.....	488
Multiple console support (MCS) in VTAM.....	489
VTAM usage of IBM Function Registry for z/OS.....	489
Controlling the domain.....	490
Establishing and terminating sessions with operator commands.....	490
Dynamic table replacement.....	491
Deactivating resources.....	491

Order of deactivation.....	491
Automatic deactivation.....	492
Normal deactivation.....	492
Immediate deactivation.....	492
Forced deactivation.....	492
Forced reactivation.....	493
Halting VTAM.....	493
Canceling VTAM.....	494
Automatic operations.....	494
Program operators.....	494
CNM application programs.....	495
Collecting session awareness (SAW) data.....	495
Chapter 20. Tuning VTAM for your environment.....	499
Introduction to tuning.....	499
Tuning tools.....	499
Estimating active sessions.....	500
Common storage areas.....	500
Buffer pools.....	500
Coattailing.....	500
Tuning tools.....	500
Monitoring common storage areas.....	500
Analyzing slowdown.....	502
Gathering tuning information with the performance monitor interface.....	502
Gathering tuning statistics.....	503
Analyzing tuning statistics.....	526
Determining the amount of coattailing in your system.....	527
Migrating from user-replaceable constants.....	527
Estimating the number of active sessions.....	528
EAS operand for application programs.....	528
EAS operand for independent logical units.....	528
Common storage areas.....	528
Common service area limit.....	528
Common service area 24-bit.....	529
DISPLAY STORUSE pools.....	529
Buffer pools.....	536
Types of buffer pools.....	536
Buffer pool allocation.....	539
HOT I/O detection/termination.....	545
Maximizing coattailing.....	546
Controlling outbound coattailing.....	547
Controlling inbound coattailing.....	547
Coattailing for SNA controllers.....	548
Coattailing for channel-to-channel operations.....	554
Session-level pacing tuning considerations.....	555
Input/output buffers.....	555
Application program data space.....	555
CSM storage.....	556
Appendix A. TSO/VTAM.....	557
Defining TSO to VTAM.....	557
Defining the TCAS application to VTAM.....	557
Defining TSO/VTAM session parameters.....	561
Defining an interpret table for compatible logons.....	566
Defining TSO to MVS.....	566
Writing a procedure to start TSO/VTAM time sharing.....	566
Creating a TSOKEY00 PARMLIB member.....	566

Defining TCAS program properties.....	566
Implementing TSO/VTAM.....	566
Translation tables.....	567
Coding TSO/VTAM exit routines.....	567
Security.....	567
Performance.....	567
3270 large screen considerations.....	568
TSO considerations.....	568
Multicultural support for TSO/VTAM user messages.....	569
Operating VTAM under TSO.....	569
Appendix B. Storage estimate worksheets.....	571
General information.....	572
APPN interchange node or network node.....	572
APPN migration data host and end node.....	575
Subarea data host.....	576
Subarea communication management configuration.....	577
APPL EAS storage estimates.....	579
Appendix C. Communications storage manager.....	581
CSM installation and definition.....	581
Initializing CSM.....	582
Monitoring CSM.....	582
CSM problem diagnosis.....	583
Appendix D. Logon manager.....	585
How the logon manager operates.....	585
Installing the logon manager.....	585
Starting the logon manager.....	586
Defining the logon manager.....	586
Sample logon manager configuration.....	586
Defining the logon manager and TPF applications to VTAM.....	588
Defining the logon manager configuration.....	588
Monitoring logon manager resources.....	592
Halting the logon manager.....	592
Appendix E. Cryptographic keys.....	593
Filing SLU keys for single-domain cryptographic sessions.....	593
Single-domain cryptographic sessions that use PCF/CUSP.....	593
Single-domain cryptographic sessions that use ICSF/MVS.....	594
Filing CDRM keys for cross-domain cryptographic sessions.....	594
Cross-domain cryptographic sessions in which both hosts use PCF/CUSP.....	595
Cross-domain cryptographic sessions in which both hosts use ICSF/MVS.....	596
Cross-domain cryptographic sessions in which the hosts use different cryptographic services....	597
Changing the cryptographic capability of a logical unit.....	599
Appendix F. Command lists: Dynamic configuration of channel-attached devices	601
ISTDINFO: VTAM device information services.....	601
Dependencies and restrictions.....	601
Output variable.....	601
Output tokens.....	602
ISTDEFIN: VTAM device information services.....	605
Dependencies and restrictions.....	605
Output variable.....	606
Output tokens.....	606

Appendix G. Message translation using the MVS Message Service.....	609
Overview of MMS support.....	609
Internal translation.....	609
Selecting internal translation.....	609
Defining messages for internal translation.....	609
External translation.....	610
Selecting external translation.....	610
Defining messages for external translation.....	610
Skeleton file use.....	611
Appendix H. Forcing an APPN route in a VTAM network.....	613
Appendix I. Border node connection types.....	617
Appendix J. VTAM restricted materials.....	621
Appendix K. Architectural specifications.....	623
Appendix L. Architectural specifications.....	625
Appendix M. Accessibility.....	627
Notices.....	629
Terms and conditions for product documentation.....	630
IBM Online Privacy Statement.....	631
Policy for unsupported hardware.....	631
Minimum supported hardware.....	631
Policy for unsupported hardware.....	632
Trademarks.....	632
Bibliography.....	633
Index.....	637

Figures

1. Pacing flow – outbound pacing.....	17
2. Pacing flow – inbound pacing.....	18
3. Pacing flow – receiving early pacing response.....	18
4. VTAM network.....	20
5. Configuration list.....	28
6. MPC connection between two VTAM network nodes.....	43
7. Type 2.1 channel connection between a composite network node and a network node.....	47
8. Leased connection between two composite network nodes.....	48
9. Two network nodes connected using an IBM 3172 Nways Interconnect Controller.....	49
10. VTAM attachment to a LAN—No meshed connection definitions.....	50
11. VTAM attachment to a LAN—Meshed connection definitions provide optimal route calculation.....	51
12. VTAM attachment to a connection network reduces required connection definitions (token ring).....	52
13. VTAM attachment to a connection network also enables optimal route calculation (token ring).....	53
14. Basic ATM configuration.....	56
15. VTAM connection to the IBM Open Systems Adapter.....	56
16. Definition of VTAM connection to the IBM Open Systems Adapter.....	57
17. Port on the IBM Open Systems Adapter through which the ATM network is accessed.....	59
18. Definition of port on the IBM Open Systems Adapter through which the ATM network is accessed.....	59
19. TGs that route data across the ATM network.....	60
20. Definition of a TG over a PVC.....	61
21. TG over an SVC.....	62
22. Definition of a TG over an SVC.....	63
23. Multiple nodes communicating across an ATM network.....	66

24. ATM configuration with a connection network.....	67
25. Definitions in VTAMLST for the VTAM in HOST1.....	69
26. Definitions in VTAMLST for the VTAM in HOST2.....	70
27. Definitions in VTAMLST for the VTAM in HOST3.....	71
28. APPN subnetworks through APPN multiple network connectivity support.....	76
29. VR-based TG between composite network nodes.....	81
30. Multiple contiguous VR-based TGs.....	82
31. VR-based TGs in a communication management configuration.....	84
32. Parallel transmission groups in multiple domain environment with NCP.....	88
33. Comparison between an EE network and a SNA network.....	93
34. How EE and EBN work together.....	94
35. How ToS settings affect IP traffic.....	98
36. Four types of EE connectivity.....	100
37. VRN connectivity.....	117
38. VTAM routing with an SATF.....	118
39. VTAM routing meshed connections.....	119
40. VTAM attachment to a connection network.....	120
41. VTAM optimal route calculations.....	121
42. Defining multiple EE VRNs.....	123
43. EE connection network reachability awareness in a mixed-release environment.....	140
44. Connection network reachability example 1.....	143
45. Connection network reachability example 2.....	143
46. Global VRN with extended border nodes.....	150
47. OSA-Express virtual MAC routing.....	162
48. QDIO primary and secondary routing.....	164

49. Creating resources in dynamic switched major node.....	168
50. Dynamic configuration of channel-attached device.....	171
51. Nonnative network type 2.1 connection.....	184
52. Local SNA devices.....	186
53. Independent LU with multiple connections to VTAM.....	194
54. Definition building for dynamically defined dependent LUs.....	198
55. Macroinstructions for logon mode table.....	205
56. How session parameters are obtained from a logon mode table.....	209
57. Automatic logon to A50ACCTS application program.....	211
58. Automatic logon reallocation.....	213
59. Fixed-session pacing (one- and two-stage).....	218
60. Adaptive session pacing.....	219
61. Pacing windows for SNA LUs.....	222
62. Pacing windows for non-SNA LUs.....	222
63. Same domain application program-to-application program session.....	225
64. Same domain application program-to-local device session.....	226
65. Application program-to-application program over APPN host-to-host connection.....	227
66. Application program-to-application program over CTCA connection.....	229
67. Application program-to-local SNA device over CTCA connection.....	230
68. Application program-to-local SNA device over AHHC connection.....	231
69. Application program-to-application program with VR from intermediate host-to-SLU host.....	232
70. Routing example through a CNN node.....	251
71. CNN route calculation example.....	252
72. Composite network node route calculation example.....	252
73. Composite network node route BIND reroute example.....	253

74. Typical CNN session path.....	254
75. APPN view of CNN session path.....	254
76. SAMAP session routing concept.....	254
77. SAMAP example.....	255
78. One explicit route in each direction.....	257
79. Two explicit routes in each direction.....	258
80. Virtual route and explicit route associations and transmission priority.....	259
81. Class of Service substitution.....	262
82. Class of Service hierarchy.....	263
83. Element and subarea address incompatibility in multiple-domain environment.....	268
84. Local flow control.....	272
85. Parallel sessions using parallel transmission groups.....	274
86. Sample single-domain network.....	275
87. Data compression yield.....	287
88. Encryption facility specifications.....	291
89. Encryption facility in an APPN environment.....	294
90. Encryption facility in multiple-network environment.....	295
91. 3270 IDS protection overview.....	300
92. Candidate application assessment process.....	303
93. Sample of typical SNA 3270 network configuration.....	305
94. Sample extended recovery facility network.....	321
95. Application program backup using persistent LU-LU sessions - part 1.....	328
96. Application program backup using persistent LU-LU sessions - part 2.....	328
97. Application program backup using persistent LU-LU sessions - part 3.....	329
98. Example of flows between client and server for DISPLAY APING command.....	338

99. Generic resource mapping.....	354
100. Partner LU mapping.....	355
101. Session establishment with generic resource members.....	356
102. Application program backup using multinode persistent sessions.....	367
103. Multinode persistent session network example.....	369
104. Session routes.....	371
105. Path switch processing.....	372
106. HPR=(RTP,ANR) and TG capabilities.....	392
107. HPR over composite network nodes.....	393
108. HPR Over APPN host-to-host channel connections.....	394
109. Multiple HPR routes between HPR session endpoints.....	395
110. Using VR-based TGs for non-HPR endpoints.....	396
111. Session involving HPR and APPN routes.....	397
112. Interchange node using HPR routing between subarea and APPN.....	397
113. Sessions traversing APPN and subarea networks.....	398
114. Example of communication management configuration.....	401
115. Communication management configuration after conversion.....	402
116. VTAM functioning as a dependent LU server.....	406
117. Switched major node for a dependent LU server.....	407
118. DLUS located in different APPN subnetwork than DLUR or PLU.....	409
119. PLU exists in or through a subarea network.....	410
120. SSCP takeover when adjacent CP does not support CP name change.....	411
121. SSCP takeover when adjacent CP is another composite network node.....	412
122. Disjoint hosts.....	415
123. Hosts with subarea connection.....	416

124. LOGMODE resolution example.....	416
125. Class of Service resolution at multiple nodes.....	419
126. Multiple-domain network.....	421
127. Major and minor nodes in multiple-domain environment.....	424
128. Example of adjacent SSCP table connection.....	433
129. Multiple-network environment.....	440
130. Multiple-network configuration: CDRM major nodes.....	446
131. Example of three interconnected networks.....	448
132. Example of two interconnected networks.....	450
133. Multiple-network configuration.....	455
134. Multiple-network paths.....	460
135. COS resolution in a multiple-network environment.....	462
136. COS tables and routing in an SNI back-to-back configuration.....	463
137. Address translation.....	464
138. Example of name translation.....	469
139. Nonnative network type 2.1 connections.....	473
140. Restoring resource definitions with configuration restart and NODELST files.....	478
141. Example of an MPC channel-to-channel, QDIO, and Hipersockets tuning statistics report.....	517
142. Buffer pool after initial allocation and after one expansion.....	540
143. How VTAM uses input buffers.....	544
144. How VTAM uses output buffers.....	544
145. Effect of DELAY time on coattailing - example 1.....	549
146. Effect of DELAY time on coattailing - example 2.....	549
147. Effect of DELAY time on coattailing - example 3.....	550
148. General I/O buffer format.....	552

149. Using multiple I/O buffers to transfer single message.....	552
150. Multiple-buffer considerations.....	553
151. Sample Logon Manager Configuration.....	587
152. Cryptography in multiple-domain environment (Both hosts use PCF/CUSP).....	596
153. Cryptography in multiple-domain environment (Both hosts use ICSF/MVS).....	597
154. Cryptography in multiple-domain environment (Hosts use different cryptographic services).....	599
155. Sample network showing default route.....	613
156. Sample network using default for UPARM1.....	614
157. Sample network using TG profile on some links.....	614

Tables

1. LU types.....	12
2. Pacing types.....	17
3. MVS static system symbols that can be used in VTAM.....	32
4. HPDT Packing - packing buffer size, CSM pools size, and waste per packing buffer.....	46
5. Comparison between static and dynamic definitions.....	99
6. SNA priorities and corresponding port numbers and default ToS values.....	136
7. Connection conditions and results.....	145
8. Troubleshooting EE problems.....	152
9. Definitions for dynamically configured devices.....	172
10. Dynamic reconfiguration operations for valid major nodes.....	175
11. Rules for multiple definition of resources.....	180
12. Example of model name table.....	202
13. Example of associated LU table.....	203
14. How session parameters are identified.....	205
15. Correspondence of methods to letters.....	221
16. Same domain application program-to-application program session—PLU to SLU flow.....	225
17. Same domain application program-to-application program Session—SLU to PLU Flow.....	226
18. Same domain application program-to-local device session—PLU to SLU flow.....	226
19. Same domain application program-to-local device session—SLU to PLU flow.....	227
20. Application program-to-application program over APPN host-to-host connection—PLU to SLU flow	228
21. Application program-to-application program over APPN host-to-host connection—SLU to PLU flow	228
22. Application program-to-application program over CTCA connection—PLU to SLU flow.....	229
23. Application program-to-application program over CTCA connection—SLU to PLU flow.....	229

24. Application program to local SNA device over CTCA connection—PLU to SLU flow.....	230
25. Application program-to-local SNA device over CTCA connection—SLU to PLU flow.....	230
26. Application program-to-local SNA device over AHHC connection—PLU to the SLU.....	231
27. Application program-to-local SNA device over AHHC connection—SLU to PLU flow.....	231
28. Application program-to-application program with VR from intermediate host-to-SLU host—PLU to SLU flow.....	232
29. Application program-to-application program with VR from intermediate host-to-SLU host—SLU to PLU flow.....	232
30. Default registration values for resources.....	240
31. COSAPPN #CONNECT class of service LINEROW values.....	244
32. ISTACST2 #CONNECT class of service LINEROW values.....	244
33. ISTACST3 #CONNECT class of service LINEROW values.....	245
34. ISTINCLM APPNCOS values based on LOGMODE operand values.....	247
35. Default TG characteristics.....	249
36. #CONNECT Class of Service LINEROW values.....	249
37. Adjacent node subarea requirements for multiple-domain environment.....	269
38. Endpoint node subarea requirements for multiple-domain environment.....	269
39. Sample model application program names.....	279
40. Compression values for example of data compression yield.....	287
41. DES-TDES24 encryption options.....	292
42. Resource verification reduction matrix.....	332
43. Structure attributes used to compute structure size.....	344
44. Structure attributes defined by VTAM.....	345
45. VTAM to VTAM connection example.....	351
46. Start options and node type relationship.....	385
47. Network resource list example.....	471
48. EQDIO Global Statistics record	505

49. EQDIO Control Read record	507
50. EQDIO Control Write record	507
51. EQDIO Datapath Read record.....	508
52. EQDIO Datapath Write record.....	510
53. Record format for SNA controller.....	511
54. Record format for channel-to-channel adapters.....	513
55. Record format for multipath channel connections (XCF).....	514
56. Record format for multipath channel connections (channel).....	517
57. Record format for TCP connections.....	522
58. Record format for RoCE connections.....	524
59. DISPLAY STORUSE pools.....	529
60. VTAM buffer pools.....	536
61. Number of buffers per page.....	542
62. I/O buffer size and number of buffers per page.....	543
63. Coding the device-specific hexadecimal data of PSERVIC.....	563
64. Worksheet for APPN interchange node or network node storage.....	572
65. Summary of worksheet, APPN interchange node or network node storage.....	574
66. Worksheet for APPN migration data host and end node	575
67. Summary of APPN migration data host and end node.....	575
68. Worksheet for subarea data host.....	576
69. Summary of subarea data host.....	576
70. Worksheet for subarea communication management configuration.....	577
71. Summary of subarea communication management configuration.....	578
72. Buffer pools in CSM.....	581
73. Connection type for selected VTAM and partner node combinations.....	617

74. User replaceable or modifiable modules.....	621
75. VTAM message modules.....	621

About this document

This document provides an understanding of the major concepts involved in implementing a VTAM® network and describes how to:

- Install VTAM
- Define network resources to VTAM
- Replace user tables and exit routines
- Tune VTAM for a specific environment

It also provides guidance for using VTAM resource definition statements and macroinstructions. Use this document when installing, upgrading, or otherwise changing a VTAM network.

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.

The [z/OS Communications Server: SNA Resource Definition Reference](#) contains a brief description and shows the exact coding format of each VTAM definition statement, start option, and user table. It serves as a reference document for system programmers who are already familiar with the major concepts involved in implementing a VTAM network.

Who should read this document

System programmers who are familiar with the basic concepts of telecommunication, Systems Network Architecture (SNA), and VTAM.

For an overview of the new functions in VTAM, see the [z/OS Communications Server: New Function Summary](#).

How this document is organized

The [z/OS Communications Server: SNA Network Implementation Guide](#) is divided into the following sections:

The first section describes subarea and Advanced Peer-to-Peer Networking (APPN) concepts for VTAM.

The second section presents an overview of VTAM implementation issues common to all VTAM networks.

The third section describes:

- Various ways of connecting VTAM as an APPN node in the network
- Enterprise Extender connections
- Types of subarea connections to VTAM
- VTAM as a data host
- Methods for VTAM resources to be defined dynamically
- Details about how to define PUs to VTAM
- Details about how to define independent and dependent LUs to VTAM

The fourth section describes:

- Major concepts involved in establishing sessions.
- Major concepts involved in routing data through a VTAM network; both routing considerations for APPN and subarea nodes are covered.

The fifth section describes:

- Concepts and functions available for VTAM application programs. For detailed information about implementation of VTAM application programs, see [z/OS Communications Server: SNA Programming](#) and [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

The sixth section describes how to incorporate VTAM functions and enhancements used in a sysplex environment.

The seventh section presents the major concepts in implementing an APPN network and describes:

- How to implement VTAM as an APPN node in the network
- Considerations for implementing a combined APPN and subarea network

The eighth section presents the major concepts involved in implementing a subarea network and describes:

- Major concepts in implementing a basic subarea network environment
- Major concepts in implementing a multiple-network environment using the SNA network interconnection (SNI)

The ninth section describes VTAM operator control functions including the basic tasks available to the VTAM operator for controlling the VTAM domain and how to control a domain that includes NCP subarea nodes.

The tenth section describes how you can adjust VTAM to provide optimal service for your environment.

The provide additional information for this document.

How to use this document

To use this document, you should be familiar with the basic concepts of telecommunication, SNA, and VTAM.

How to provide feedback to IBM

We welcome any feedback that you have, including comments on the clarity, accuracy, or completeness of the information. See, [How to send feedback to IBM®](#) for additional information.

Conventions and terminology that are used in this information

Commands in this information 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 information are *installation-wide exit routines*. The installation-wide exit routines also called installation-wide exits, exit routines, and exits throughout this information.

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 information might not be updated for each release. Evaluate a sample carefully before applying it to your system.

z/OS no longer supports mounting HFS data sets (The POSIX style file system). Instead, a z/OS File System (zFS) can be implemented. The term hierarchical file system, abbreviated as HFS, is defined as a data structure that has a hierarchical nature with directories and files. References to hierarchical file systems or HFS might still be in use in z/OS Communications Server publications.

Network Express and Open Systems Adapter-Express (OSA-Express) terminology:

- The Network Express feature is introduced with the IBM z17 processor family. The Network Express feature is the next generation of Open Systems Adapter (OSA) technology. The term OSA (Open Systems Adapter) is carried forward with Network Express. The IBM z17 processor supports both the Network Express and the OSA-Express7S features. In this information, when a general reference is made to OSA that applies to all these features, then the term OSA is used, and the acronym will appear in italics. This formatting style and guideline for usage for the term OSA is used throughout this document. When a distinction is necessary, then the specific feature name is used such as the Network Express feature
- The Network Express feature is defined as channel (CHPID) type OSH (Open System Adapter for Hybrid networks) that might operate in either 10 GbE or 25 GbE link speed. When this term is used in this information, the processing being described applies to either link speed. If processing is applicable to only one link speed, the full terminology, for instance, IBM 25 GbE Network Express will be used.
- Network Express is defined with new system architecture called Enhanced Queued Direct I/O (EQDIO). In this information there are many references to QDIO or OSA/QDIO. When the reference applies to both QDIO and EQDIO the reference just indicates OSA. When the reference is specific to the QDIO or EQDIO architecture, then the specific architecture is referenced, for example, OSA/QDIO or OSA/EQDIO. Some OSA references also use or include the channel type for OSA such as OSD (QDIO). When the reference applies to both features, then the term OSA is used. When a distinction is necessary then the specific channel or architecture type is used, OSD/QDIO or OSH/EQDIO.

Shared Memory Communications over Remote Direct Memory Access (SMC-R) terminology

- *RoCE* , which is a generic term representing IBM® 10 GbE RoCE Express, IBM 10 GbE RoCE Express2, IBM 25 GbE RoCE Express2, IBM 10 GbE RoCE Express3, IBM 25 GbE RoCE Express3, IBM 10 GbE Network Express and IBM 25 GbE Network Express feature capabilities. When this term is used in this information, the processing being described applies to all of these features. If processing is applicable to only one feature, the full terminology, for instance, Network Express will be used.
- RoCE Express2, which is a generic term representing an IBM RoCE Express2 feature that might operate in either 10 GbE or 25 GbE link speed. When this term is used in this information, the processing being described applies to either link speed. If processing applies to only one link speed, the full terminology, for instance, IBM 25 GbE RoCE Express2 will be used.
- RoCE Express3, which is a generic term representing an IBM RoCE Express3 feature that might operate in either 10 GbE or 25 GbE link speed. When this term is used in this information, the processing being described applies to either link speed. If processing applies to only one link speed, the full terminology, for instance, IBM 25 GbE RoCE Express3 will be used.
- Network Express, which is a generic term representing an Network Express feature that might operate in either 10 GbE or 25 GbE link speed. When this term is used in this information, the processing being described applies to either link speed. If processing is applicable to only one link speed, the full terminology, for instance, IBM 25 GbE Network Express will be used. When configured with a CHPID type of NETH, the Network Express feature may operate as an RDMA network interface card.
- RDMA network interface card (RNIC), which is used to refer to the IBM 10 GbE RoCE Express, IBM 10 GbE RoCE Express2, IBM 25 GbE RoCE Express2, IBM 10 GbE RoCE Express3, or IBM 25 GbE RoCE Express3, IBM 10 GbE Network Express or IBM 25 GbE Network Express feature.
- Shared RoCE environment, which means that the *RoCE* feature can be used concurrently, or shared, by multiple operating system instances. The feature is considered to operate in a shared RoCE environment even if you use it with a single operating system instance.

Clarification of notes

Information traditionally qualified as Notes is further qualified as follows:

Attention

Indicate the possibility of damage

Guideline

Customary way to perform a procedure

Note

Supplemental detail

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


Tip

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


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).

►► 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)**.

►► 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)**.

►► 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.

►► *variable* ◄◄

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

►► OPERAND — ¹ ◄◄

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.

►► Reference to Syntax Fragment ◄◄

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 633, 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	Description
<i>IBM Z Redbooks</i>	The IBM Z® subject areas range from e-business application development and enablement to hardware, networking, Linux®, solutions, security, parallel sysplex, and many others. For more information about the Redbooks® publications, see http://www.redbooks.ibm.com/ and http://www.ibm.com/systems/z/os/zos/zfavorites/ .

Other documents

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

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

To find the complete z/OS library, visit the [z/OS library](https://www.ibm.com/docs/en/zos) in [IBM Documentation](https://www.ibm.com/docs/en/zos) (<https://www.ibm.com/docs/en/zos>).

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

Title	Number
<i>Understanding LDAP</i>	SG24-4986
<i>z/OS Cryptographic Services System SSL Programming</i>	SC14-7495
<i>z/OS IBM Tivoli Directory Server Administration and Use for z/OS</i>	SC23-6788
<i>z/OS JES2 Initialization and Tuning Guide</i>	SA32-0991
<i>z/OS Problem Management</i>	SC23-6844
<i>z/OS MVS Diagnosis: Reference</i>	GA32-0904
<i>z/OS MVS Diagnosis: Tools and Service Aids</i>	GA32-0905
<i>z/OS MVS Using the Subsystem Interface</i>	SA38-0679
<i>z/OS Program Directory</i>	GI11-9848
<i>z/OS UNIX System Services Command Reference</i>	SA23-2280
<i>z/OS UNIX System Services Planning</i>	GA32-0884
<i>z/OS UNIX System Services Programming: Assembler Callable Services Reference</i>	SA23-2281
<i>z/OS UNIX System Services User's Guide</i>	SA23-2279
<i>z/OS C/C++ Runtime Library Reference</i>	SC14-7314
<i>OSA-Express Customer's Guide and Reference</i>	SA22-7935

Redbooks publications

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

Title	Number
<i>IBM z/OS Communications Server TCP/IP Implementation, Volume 1: Base Functions, Connectivity, and Routing</i>	SG24-8096
<i>IBM z/OS Communications Server TCP/IP Implementation, Volume 2: Standard Applications</i>	SG24-8097
<i>IBM z/OS Communications Server TCP/IP Implementation, Volume 3: High Availability, Scalability, and Performance</i>	SG24-8098
<i>IBM z/OS Communications Server TCP/IP Implementation, Volume 4: Security and Policy-Based Networking</i>	SG24-8099
<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

<http://www.ibm.com/systems/z/os/zos/library/bkserv/>

z/OS Communications Server product

The page contains z/OS Communications Server product introduction

<https://www.ibm.com/products/zos-communications-server>

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

<https://www.ibm.com/mysupport>

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/>

z/OS Support Community

Search the z/OS Support Community Library for Techdocs (including Flashes, presentations, Technotes, FAQs, white papers, Customer Support Plans, and Skills Transfer information)

[z/OS Support Community](#)

Tivoli® NetView for z/OS

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

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

RFCs

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

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

Internet drafts

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

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

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

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

DNS websites

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

USENET news groups

comp.protocols.dns.bind

BIND mailing lists

<https://lists.isc.org/mailman/listinfo>

BIND Users

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

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

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

The z/OS Basic Skills Information Center

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

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

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

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

Summary of changes for SNA Network Implementation Guide

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

Summary of changes for z/OS 3.2

The following content is new, changed, or no longer included in z/OS 3.2.

Changed

The following content is changed.

September 2025 release

- None.

Changes made in z/OS Communications Server 3.1

The following content is new, changed, or no longer included in z/OS 3.1.

New information

- Added new information for EQDIO, see
 - [“Tuning input/output \(I/O\) operations for EQDIO interfaces” on page 504](#)
- Communications Server exploitation of the IBM Function Registry for z/OS, see the following topic:
 - [“VTAM usage of IBM Function Registry for z/OS” on page 489](#)

Changed information

- Information for REUSASID is added to [“Recommended start option” on page 22](#).
- Communications Server support for RoCE Express3, see the following topic:
 - [“Resources automatically activated by VTAM” on page 483](#)
- Removal of OSA DEVICE/LINK/Home Configuration Support, see the following topics:
 - [“Using MPC connections in an APPN network” on page 41](#)
 - [“OSA-Express overview” on page 159](#)
 - [“Defining an OSA-Express device to z/OS Communications Server using QDIO” on page 159](#)
 - [“OSA-Express virtual MAC \(VMAC\) routing ” on page 161](#)
 - [“Primary and secondary routing” on page 163](#)
 - [“Displaying and modifying 3270 IDS configuration” on page 309](#)
 - [“Resources automatically activated by VTAM” on page 483](#)
- Withdrawal of support of VTAM LSA and TCP/IP LCS devices, see the following topics:
 - [“Routing in a subarea network” on page 13](#)
 - [“Using MPC connections in an APPN network” on page 41](#)
 - [“HPDT Packing” on page 45](#)

- [“Overriding dial number digits for dial connections” on page 283](#)
- [“OSA-Express overview” on page 159](#)
- [“Defining an OSA-Express device to z/OS Communications Server using QDIO” on page 159](#)
- [“Primary and secondary routing” on page 163](#)
- [“Dynamic definition of independent LUs” on page 191](#)
- [“Independent logical units and adaptive-session pacing” on page 223](#)
- [“What are the IBM-supplied default classes of service?” on page 243](#)
- [“Physical paths” on page 256](#)
- [“What is Rapid Transport Protocol?” on page 389](#)
- [“Limiting HPR support” on page 391](#)
- [“Resources automatically activated by VTAM” on page 483](#)
- [“Tuning statistics file report for TCP connections” on page 522](#)
- [“Guidelines for setting DELAY” on page 549](#)

Deleted information

- Withdrawal of support of VTAM LSA and TCP/IP LCS devices. The following topics have been deleted:
 - *IBM 3172 Nways Interconnect® Controller connections between APPN nodes*
 - *Examples*
 - *Using a connection network*
 - *IBM Open Systems Adapter connections between APPN nodes*
 - *Connecting two VTAMs using an external communication adapter*
 - *Sample configuration with Ethernet or Ethernet-type LAN*
 - *Sample configuration with a token-ring local area network*
 - *External communication adapter (XCA) connections*
 - *Operating VTAM in a multiple-domain subarea network*

Chapter 1. Post-installation considerations for z/OS Communications Server

This section describes some post-installation considerations for Communications Server under the z/OS operating system.

Defining Communications Server SNA to z/OS

To define z/OS Communications Server under the z/OS operating system, you need to change the message routing codes, define channel-attached devices and determine the ECSA value.

Procedure

To define z/OS Communications Server, do the following steps:

1. Change the message routing codes by coding a system user exit routine (if there is multiple console support).

Message routing codes determine the console at which messages will appear. If the routing codes provided for z/OS Communications Server messages do not meet your needs, you can change the routing codes used on the messages by coding a system user exit routine (if there is multiple console support). The exit routine receives control before messages are routed so it can examine the messages' routing codes (and descriptor codes) and change them. The system uses the modified routing codes to route these messages. To change routing codes, do the following steps:

- a. Prepare the write-to-operator/write-to-operator-with-response (WTO/WTOR) exit routine, and add it to the control program. The WTO/WTOR exit routine can be inserted into the resident portion (communications task) of the control program either before or after system generation.
 - b. See [z/OS Communications Server: SNA Messages](#) for the message routing codes, and decide which new routing codes you want to assign to each message.
2. Define channel-attached devices.

If you are adding channel-attached devices, you can define these devices using the Hardware Configuration Definition (HCD) to dynamically add the devices.

3. Determine the ECSA value.

The ECSA value is defined by the second value on the CSA parameter in member IEASYSxx. It is recommended that you examine the ECSA value to ensure that it is adequate for z/OS Communications Server. You can determine the ECSA value by reviewing [Appendix B, "Storage estimate worksheets,"](#) on page 571.

Using automatic restart manager

Automatic restart manager is a z/OS function that can automatically restart z/OS Communications Server after an abnormal end (abend).

During initialization, z/OS Communications Server automatically registers with the automatic restart manager, using the following options:

- REQUEST=REGISTER
- ELEMENT=NET@cp_name
- EVENTEXIT=NO_EVENTEXIT
- STARTTXT=NO_STARTTXT
- ELEMTYPE=SYSVTAM
- TERMTYPE=ELEMTERM

Note: The *cp_name* is the same name as that used on the SSCPNAME start option.

For more information about automatic restart manager, see [z/OS MVS Setting Up a Sysplex](#).

Starting z/OS Communications Server SNA

You should code a z/OS Communications Server start procedure and save it in SYS1.PROCLIB. The system operator specifies the procedure when starting z/OS Communications Server.

The start procedure is called NET. The name NET is not required but is strongly recommended for consistency in entering the z/OS Communications Server operator commands and to reduce the operator's chances of making a syntax error. The procedure name you specify must be the first operand on the START, and MODIFY operator commands. For DISPLAY, HALT, and VARY the procedure name is always NET.

Following is an example of job control statements for a typical start procedure.

```
//NET      PROC
//VTAM     EXEC PGM=ISTINM01,TIME=1440,REGION=4096K,
//STEPLIB DD DSN=SYS1.SSPLIB,DISP=SHR
//FFSTLIB DD DSN=SYS1.VTAMLIB,DISP=SHR
//VTAMLST DD DSN=SYS1.VTAMLST,DISP=SHR
//         DD DSN=USER1.AUTO.VTAMLST,DISP=SHR
//SISTCLIB DD DSN=SYS1.SISTCLIB,DISP=SHR
//AUTOCKPT DD DSN=VSAM.AUTOCKPT,AMP=AMORG,DISP=OLD
//VTAMLIB DD DSN=SYS1.VTAMLIB,DISP=SHR
//NCPDUMP DD DSN=SYS1.NCPDUMP,DISP=SHR
//* DATA SETS FOR APPN DATABASE CHECKPOINTING
//DSDB1 DD DSN=SYS1.DSDB1,DISP=SHR
//DSDB2 DD DSN=SYS1.DSDB2,DISP=SHR
//DSDBCTRL DD DSN=SYS1.DSDBCTRL,DISP=SHR
//TRSDDB DD DSN=SYS1.TRSDDB,DISP=SHR
//* DATA SETS 3720, 3725, AND 3745 DUMPS
//LDRIOTAB DD DSN=SYS1.LDRIOTAB,DISP=(SHR,PASS,KEEP)
//CSPDUMP DD DSN=SYS1.CDUMP,DISP=SHR
//MOSSDUMP DD DSN=SYS1.MDUMP,DISP=SHR
//NCPLOAD DD DSN=SYS1.NCPLOAD,DISP=SHR
//* NODELST DATA SET
//NODLST1 DD DSN=VSAM.NODLST1,AMP=AMORG,DISP=OLD
//* ALTERNATE NODELST DATA SET
//NODLST2 DD DSN=VSAM.NODLST2,AMP=AMORG,DISP=OLD
```

Note: On the EXEC statement, ISTINM01 is the main z/OS Communications Server initialization module name. Code PGM=ISTINM01.

The previous example is based on the following assumptions:

- The node is APPN capable.
- A communication controller is in the network.
- The generated NCP and RRT modules for the communication controller reside in data set SYS1.NCPLOAD (NCPLOAD DD statement). The NCP source is in SYS1.VTAMLST.
- The SSP modules needed to load and dump the communication controllers are in SYS1.SSPLIB (STEPLIB DD statement).
- A dump data set is needed for the communication controller (NCPDUMP DD statement).
- The dynamic configuration of channel-attached devices facility (dynamic I/O) is being used (VTAMLST DD and AUTOCKPT DD statements).
- A dump data set is needed for the communication scanner processor (CSP) (CSPDUMP DD statement). For 3380 DASD, this data set should be allocated with at least 7 cylinders using a block size of 512.
- A dump data set is needed for the maintenance and operator subsystem (MOSS) (MOSSDUMP DD statement). For 3380 DASD, this data set should be allocated with at least 10 cylinders using a block size of 512.
- Two NODELST data sets have been defined and can be used by z/OS Communications Server (NODEDS1 DD and NODEDS2 DD statement).
- Space has been allocated for all data sets, and they have been cataloged.

Chapter 2. VTAM networking concepts

This topic describes Virtual Telecommunications Access Method (VTAM) concepts for a subarea and Advanced Peer-to-Peer Networking (APPN) environments. VTAM is an application program that is based on System Network Architecture (SNA). This topic includes:

- VTAM functions
- VTAM nodes
- Network accessible units
- Network control sessions
- User sessions
- How VTAM locates resources
- Route selection
- Addressing
- Controlling network data flow using pacing

VTAM functions

VTAM performs a number of tasks in a network. For example, VTAM:

- Monitors and controls the activation and connection of resources
- Establishes connections and manages the flow and pacing of sessions
- Provides application programming interfaces (for example, an APPC API for LU 6.2 programming) that allow access to the network by user-written application programs and IBM-provided subsystems
- Provides interactive terminal support for Time Sharing Option (TSO) using Multiple Virtual Storage (MVS)
- Provides support for both locally and remotely attached resources

VTAM nodes

VTAM can be configured as one of the following types of nodes:

- APPN function only
- Subarea function only
- APPN and subarea function

The type you choose depends on the level of function required for the node and whether the node needs to communicate with other subarea nodes, other APPN nodes, or both.

Nodes with APPN function only

A VTAM APPN node is an SNA type 2.1 node that functions in a peer-to-peer environment. APPN nodes use APPN or low entry network (LEN) connections to communicate with other nodes in an APPN network.

A VTAM APPN node:

- Does not require subarea network routing definitions
- Uses control point-control point (CP-CP) sessions for network control data and logical unit-logical unit (LU-LU) sessions for user data
- Does not support system services control program-system services control program (SSCP-SSCP) sessions
- Has the NODETYPE start option defined

- Does not have a subarea number defined on the HOSTSA start option

Note: If the start option, SACONNS, is set to NO, VTAM HOSTA is a pure APPN node even if HOSTA is coded.

- Cannot activate a network control program (NCP)
- Requires static definitions only for those resources located within the node
- Also provides dynamic definitions of some local connections

Control point

Each APPN node has a control point (CP) that manages the node and its resources. Control points provide function for the APPN network similar to the function the SSCP provides for the subarea network. The CP can implement end node or network node functions.

The CP:

- Activates and deactivates network resources
- Assists in session initiation and termination
- Routes cross-domain requests
- Maintains a dynamic representation of the network in the topology and directory services databases (network nodes only)

In an end node, the CP communicates only with the CP in the network node that provides it with network services. In a network node, the CP communicates with the CPs in adjacent network nodes and with the CPs in adjacent end nodes for which the CP is the network node server. The CP directs the activation and deactivation of resources and links and helps LUs start and stop sessions.

The control point name of an adjacent APPN node can be statically defined or determined during XID exchange. The control point name of a VTAM node is determined by the SSCPNAME start option.

Network node

A network node (NN) supports its own users and the end nodes it serves by providing directory and route selection services. It performs intermediate routing of data for sessions that cross it. The NN performs searches of the network to locate resources and calculates the best session route from the node of the primary LU to the node of the secondary LU, based on user-specified criteria.

A VTAM NN:

- Provides intermediate session routing
- Maintains CP-CP sessions with all adjacent network nodes and with all end nodes for which it provides network services
- Has the NODETYPE start option set to NN (NODETYPE=NN)
- Does not have a subarea number defined on the HOSTSA start option

Note: If the start option, SACONNS, is set to NO, VTAM HOSTA is a pure APPN node even if HOSTA is coded.

- Participates in searches for network resources
- Cannot activate an NCP
- Provides dynamic route calculation

Network nodes and their interconnections form an intermediate routing network. Network nodes use this intermediate routing network to perform searches for partner LUs, and to calculate the best available route for a session for their own resources and the resources of end nodes for which they are acting as network node servers.

Network nodes are distinguished from end nodes by the functions they provide. An NN provides network services for its own resources and the resources of the end nodes for which it acts as a network node server. For more information about network node servers, see [“Network node server” on page 5](#).

To perform NN-provided network services, network nodes maintain a directory services database and a topology database. The directory services database contains information about resources in the network, including network node local resources, and is updated dynamically as a result of searches for partner LUs. The topology database includes information about network nodes and transmission groups (TGs) and is updated dynamically whenever there is a change to the intermediate routing network.

All directory requests from resources (such as LUs on served end nodes, or the network node resources) use the directory services of the NN. Therefore, the network nodes are able to collect and control directory information about the APPN network. Network nodes can originate and participate in APPN broadcast searches.

End node

An end node (EN) is an APPN node that relies on the directory and route selection services of a network node to participate in an APPN network.

A VTAM EN:

- Interacts with a network node to request and receive directory and route selection services
- Supports CP-CP sessions between the end node and its network node server
- Can specify a preferred network node server by means of the NNSPREF start option
- Can define a network node server list
- Submits search requests to its network node server and obtains the session route from its server when initiating LU-LU sessions
- Dynamically registers its resources with its network node server so that the resources can be found during broadcast searches
- Has the NODETYPE start option set to EN (NODETYPE=EN)
- Cannot activate an NCP
- Does not have a subarea number defined on the HOSTSA start option

Note: If the start option, SACONNS, is set to NO, VTAM HOSTA is a pure APPN node even if HOSTA is coded.

An EN provides communication services for those resources defined within the node. An EN provides directory and session services for its own LUs only.

For directory and routing services to other nodes, an end node uses the services of its network node server. An end node can have links to several network nodes for connectivity and routing, but only one of these network nodes can be used as the network node server at any given time. A VTAM EN can register its resources in the directory services database of its network node server. This allows searches from other nodes to locate the resources on this EN because VTAM end nodes are not searched. An EN does not perform intermediate routing of sessions; it can only be the origin or destination of a session.

Both a data host and an EN primarily process applications and do not provide network services for the resources attached to them. If your data host has APPN connections and does not require subarea connections, you can migrate it to an EN. (If your data host requires subarea connections, you can migrate it to a migration data host, which is described in [“Migration data host” on page 8.](#))

Network node server

A network node server (NNS) is a network node that provides resource location and route selection services to the LUs it serves. These LUs can be on the network node itself, or on end nodes that have established CP-CP sessions with this network node as their network node server.

To support the LUs on served end nodes, a network node server uses the CP-CP session to provide network information for session setup. Network node servers perform searches for partner LUs and calculate the best route to the partner LU.

Any network node can be a network node server for end nodes that are attached to it. The served end nodes are defined as being in that network node server domain.

Note: A domain in the subarea network is the set of resources controlled by a particular SSCP. A domain in the APPN network is the set of resources served by a particular network node or end node.

Central directory server

A central directory server (CDS) is a network node that builds and maintains a directory of resources throughout the network. The CDS server provides some similar function to that of a CMC host in a subarea network. While a CMC is a central place for network management information that operates and controls the network, a CDS manages only the directory functions, not the operational and network management functions that a CMC also provides. The purpose of a CDS is to reduce the number of network broadcast searches to no more than one per resource. The CDS is capable of maintaining its directory even when VTAM is halted and restarted.

VTAM network nodes and end nodes can register their resources with a central directory server, which acts as a focal point for resource location information. When a VTAM end node registers its resources, it can request that resources be registered either to its network node server only, or to both its network node server and the central directory server. Entries in a directory database can be registered, defined, or dynamic.

When a network node receives a search request for a resource for which it does not have location information, the network node first sends a directed search request to a central directory server, if there is one. The central directory server searches for information about the location of the resource in its directory. If it does not find the resource's location, the central directory server searches end nodes in its domain, other central directory servers, and, if necessary, the entire network.

If the resource is still not found, the central directory server notifies the node that originally requested the search that the search is unsuccessful.

Border node

A network node capable of APPN multiple network connectivity, also known as a border node, can maintain CP-CP connectivity with a network node with a different NETID. APPN topology information does not cross the border node connection or APPN subnetwork boundary, but search requests can, and an LU-LU session can be set up. An APPN subnetwork boundary is assumed by VTAM when a border node is connected to a network node (or border node) with a different network identifier.

An APPN subnetwork boundary can also be explicitly defined through system definition. This way, border nodes can be used to create subnetwork boundaries to adjacent network nodes (or border nodes) with the same network identifier.

For further information about border nodes, see [“APPN multiple network connectivity” on page 75](#).

Nodes with subarea function only

A VTAM subarea node is an SNA type 5 node that functions in a hierarchical environment. Subarea nodes provide services for and control over peripheral nodes. Peripheral nodes require the services of a VTAM subarea node to communicate with other peripheral nodes and subarea nodes in the network. Peripheral nodes include SNA type 2.1, 2.0 or type 1 nodes and function as distributed processors, cluster controllers, or workstations. Type 2.1 nodes can use low-entry networking (LEN) or Advanced Peer-to-Peer Networking (APPN) connections to communicate with adjacent type 2.1 nodes; however, when attaching to a VTAM subarea node, they are restricted to using LEN connections.

A VTAM subarea node:

- Requires subarea network routing definitions, such as path, virtual route (VR), and explicit route (ER) definitions to communicate with other subarea nodes
- Uses SSCP-SSCP, SSCP-PU, SSCP-LU, and LU-LU sessions for network control and user data
- Does not support CP-CP sessions
- Has a subarea number defined on the HOSTSA start option
- Does not have the NODETYPE start option defined

- Can activate NCPs and have SSCP-SSCP sessions with other VTAM subarea, interchange, and migration data host nodes

System services control point

A system services control point (SSCP) provides resource management and other session services for users in a subarea network.

To control and provide services for its subordinate nodes, an SSCP establishes sessions with the physical units (PUs) and LUs in those nodes. For example, before an LU-LU session can be established, the SSCP must use a directory of network resources to locate the partner LU. When the partner LU is located, the SSCP establishes an SSCP-LU session with that partner LU.

An SSCP provides the following functions:

- Manages resources in a subarea network in accordance with the commands that network operators issue
- Coordinates the initiation and termination of sessions between LUs in separate nodes within its domain or across domains in cooperation with other SSCPs
- Coordinates the testing and status monitoring of resources within its domain
- Performs recovery when communication fails between network components

Communication management configuration and data hosts

A communication management configuration (CMC) simplifies network management. One VTAM in the network acts as the CMC host and owns all resources except applications, which are predominately owned by other hosts called data hosts. With this configuration the entire network can be controlled by a network operator at the CMC.

A CMC configuration can also be used to enhance routing. The CMC handles session establishment and termination. The data hosts then spend more time doing application processing and less time doing session establishment and routing.

Nodes with both subarea and APPN function

Nodes that need to be able to directly connect to both subarea and APPN networks require both subarea and APPN function.

Interchange node

An interchange node resides on the border of an APPN network and a subarea network. It provides protocol conversion between subarea and APPN networks to enable the integration of the two types of networks. Because an interchange node can convert session requests from one protocol to the other and can provide intermediate routing, it can establish sessions from one type of network to the other.

An interchange node combines the function of a subarea node and a network node. It controls resources and functions as a network node in the APPN network and as an SSCP and a cross-domain resource manager (CDRM) in the subarea network. All of the characteristics described for network nodes and subarea nodes apply to interchange nodes.

An interchange node:

- Uses subarea path definitions to determine routes within the subarea network
- Uses the topology database to determine routes within APPN networks
- Uses both SSCP-SSCP and CP-CP sessions to communicate with other nodes
- Has a subarea number and NODETYPE=NN defined
- Can own and activate NCPs

The interchange node communicates network control data by using SSCP-SSCP sessions with other subarea nodes and CP-CP sessions with other APPN nodes. To enable it to participate in the subarea

network, it is defined with a unique subarea number and requires subarea path definition statements. It can be connected to other APPN nodes, LEN nodes, and subarea nodes.

Note: The SACONNS start option must be set correctly to allow VTAM to function as an interchange node.

Migration data host

A migration data host resides on the periphery of a combined APPN and subarea network. A migration data host combines the function of an end node with the function and role of a subarea data host.

A migration data host:

- Uses subarea network routing definitions
- Should not perform intermediate routing in APPN or subarea networks
- Uses CP-CP and SSCP-SSCP sessions to communicate with other nodes
- Has NODETYPE=EN defined
- Cannot activate an NCP
- Has a subarea number defined on the HOSTSA start option

Note: The SACONNS start option must be set correctly to allow VTAM to function as a migration data host.

Like a data host in a subarea network, a migration data host is dedicated to processing application programs and does not control network resources. It also participates as a cross-domain resource manager (CDRM) in the subarea network. The migration data host also functions as an end node in the APPN network. All of the characteristics previously described for end nodes apply to migration data hosts.

To enable the migration data host to participate in the subarea network, it is defined with a unique subarea number and supports subarea path definition statements.

The migration data host communicates network control data by using SSCP-SSCP sessions with other subarea nodes and CP-CP sessions with its network node server. It can be connected to other APPN nodes, LEN nodes, and subarea nodes. A migration data host can perform intermediate routing in either a subarea network or an APPN network, but it is not recommended. However, its locally attached LEN nodes and dependent LUs can establish sessions with resources in other APPN or subarea domains.

Composite network node

A composite network node (CNN) is composed of a VTAM interchange node and any NCPs that it owns. In an APPN network, it functions as a network node and appears to the APPN network as a single node. A composite network node is defined by coding the HOSTSA start option, the NODETYPE start option as NN, and by activating an NCP from that VTAM.

The composite network node can be attached to other APPN nodes and also to subarea nodes. In a composite network node where all external connections are APPN connections, the HOSTSA value is not used in communication outside the node. If the composite node is connected to other nodes by subarea connections, the HOSTSA value is used in communication with the subarea network.

Notes:

1. If the composite network node has APPN connections through its NCP, the NCP needs to be at Version 6 Release 2 or later.
2. The SACONNS start option must be set correctly to allow VTAM to function as a composite network node.

Low-entry networking node

A low-entry networking (LEN) node is a type 2.1 node that can function in both a subarea and an APPN environment. VTAM and any NCPs that it owns can present the image of a composite LEN node. It appears as a single node to all LEN or APPN nodes to which it is attached.

A LEN node provides peer-to-peer connectivity to subarea nodes, end nodes, or network nodes that are providing a LEN appearance. Unlike end nodes, the LEN node cannot establish CP-CP sessions with a network node. Therefore, a LEN node cannot automatically register its resources with a network node server.

A VTAM acting as a LEN node cannot establish SSCP-SSCP sessions across T2.1 connections with a VTAM subarea node. A LEN node can attach to a subarea network as a peripheral node. In a subarea network, LEN nodes can communicate with other LEN nodes without requiring the services of an SSCP.

In an APPN network, you can predefine the CP name of a network node adjacent to a LEN node, which enables the LEN node to send a BIND to the adjacent node. The adjacent node determines the actual location of the LU, calculates the route to it, and forwards the BIND.

Network accessible units

Network accessible units (NAUs) are elements in the network with which sessions can be established and queries can be directed. In subarea networks, these units have both a subarea number and an element address and are also called network addressable units. Although subarea addresses are not used in APPN networks, the term network accessible units is used to see these units in both types of networks.

Physical unit

Physical units (PUs) exist in type 4, type 5, type 2.0, and type 1 nodes. Combined type 2.0/2.1 nodes, such as a 3174 controller, have both a CP and PU.

A PU provides the following functions:

- Receives and acts upon requests from CPs, such as activating and deactivating links to adjacent nodes
- Manages links and link stations while accounting for the unique aspects of different link types
- Manages network addresses, virtual routes, and explicit routes in type 4 and type 5 nodes

The PU supports sessions with the SSCP in type 5 nodes.

Logical unit

A logical unit (LU) acts as the intermediary between the user and the network. Before two users can communicate, their LUs must be in session with one another. Communication occurs only between LUs of the same LU type. The LU-LU session manages the flow of data between the users.

VTAM supports LU types 0, 1, 2, 3, and 6.2. These types support the communication requirements of a variety of users. Each type supports a specific set of communication requirements. The number of the LU type is not related to the PU or node type.

LU types 0, 1, 2, and 3 support communication between application programs and peripheral equipment, such as workstations or printers. LU type 6.2 supports communication between two programs located at type 5 subarea nodes, type 2.1 peripheral nodes, or both, and between programs and devices. LU type 6.2 supports multiple concurrent sessions, called parallel sessions.

Dependent LU support

Dependent LUs are supported by VTAM in both subarea and APPN networks.

A dependent LU requires assistance from a system services control point (SSCP) to activate an LU-LU session; therefore, it requires an SSCP-LU session. Both subarea and APPN VTAM nodes provide SSCP support for dependent LUs in their domains. A VTAM APPN node provides SSCP support for those LUs in its domain even though it has no SSCP-SSCP sessions with other nodes.

Dependent LUs can use APPN and subarea connections to access VTAM application programs in other domains. For example, a dependent LU type 2 terminal session originating in one VTAM can be rerouted to CICS® on a VTAM end node.

To resources in the APPN network, a dependent LU appears to be located on the node that is providing SSCP services to the LU. Through the use of DLUS/DLUR, VTAM can establish SSCP-LU sessions across APPN links. Therefore, a VTAM node can be the owner of a dependent LU that is on an attached non-VTAM APPN node, such as an AS/400, PS/2, or 3174 controller, regardless of whether the VTAM node has a CP-CP session with the other node.

How dependent LUs are defined

Dependent LUs are owned by a VTAM node and are defined to VTAM. They are considered to be in the domain of the VTAM node. If a type 2.1 node has dependent LUs attached to it, the dependent LU must also be defined to VTAM.

Dependent LUs such as terminals and printers appear as cross-domain resources to other VTAM hosts. If the LUs are registered, and if all nodes in the network are connected by CP-CP sessions, the cross-domain resource definition statements in other hosts might not be required. Some other hosts do support the dynamic creation of CDRSCs, even in a subarea environment. In an APPN environment, other hosts will still build CDRSCs when sessions are established.

How dependent LUs participate in sessions

Dependent LUs can start a session when the primary LU is owned by VTAM. Dependent LUs can be the secondary LU in sessions with partner LUs in the APPN or subarea network.

Sessions initiated by secondary LUs (such as LU type 2 displays) to a primary LU that is in a VTAM node in the APPN network are supported. Also supported are APPN LUs that initiate sessions with application programs in the subarea network.

Independent LU support

Independent LUs do not require the services of an SSCP to initiate an LU-LU session. Therefore, they do not have SSCP-LU sessions. Independent LUs can use the services of their local control point to locate and determine a session route to a required session partner.

Independent LUs must be defined at the owning CP. At the owning CP, they are defined as applications or device-type LUs. Independent LUs can be statically defined at non-owning CPs as cross-domain resources (CDRSCs), but this is unnecessary because these resources can be defined dynamically as they are needed. The one exception is when the LU is owned by a LEN node. Because a LEN node cannot register its LUs, these LUs cannot be found as destination LUs if they are not defined in the adjacent VTAM as CDRSCs with an ALSST or as LOCADDR=0 LUs under a PU definition.

Network control sessions

Network control sessions allow two nodes in a network to communicate, enabling user sessions between logical units located at those nodes. In subarea networks, network control sessions are called SSCP-SSCP, SSCP-NCP, SSCP-PU, and SSCP-LU sessions. In APPN networks, network control sessions are called CP-CP sessions.

SSCP-SSCP sessions

SSCP-SSCP sessions provide control and coordinate session setup for cross-domain and cross-network sessions. These sessions are established between VTAMs only. SSCPs exchange capabilities during CDRM activation. SSCP-SSCP sessions are established between all hosts in a network. SSCP-SSCP sessions can be established only between two hosts with subarea function. Within one subarea network, for an LU-LU session to be established, the SSCPs that own the LUs must be in session with one another.

For dependent LUs, the SSCPs must also establish SSCP-PU and SSCP-LU sessions before an LU-LU session can be established. An SSCP establishes sessions with PUs and LUs defined to it during network activation, or, in switched links, during link activation. SSCP-LU sessions need to be established for dependent LUs only. Independent LUs can function without an SSCP-LU session.

CP-CP sessions

To perform directory services and topology and route selection services, adjacent APPN nodes throughout the APPN network use a set of two CP-CP sessions to exchange network information. CP-CP sessions are established between type 2.1 nodes. A network node or composite network node can establish CP-CP sessions with any network node or composite network node to which it has an APPN link that supports CP-CP sessions, with the exception that CP-CP session between network nodes having different Net IDs are permitted only if at least one of the network nodes supports the border node function. It also establishes CP-CP sessions with each of the end nodes that it serves. End nodes can be attached to several network nodes but can establish a CP-CP session pair with only one network node server at a time.

When two nodes have one or more CP-CP capable links between them, VTAM requires a CP-CP session to be active between the nodes, as long as none of the following conditions disallowing CP-CP sessions are true:

- The CP-CP capable links exist between two adjacent network nodes with different net IDs, and neither network node provides border node support.
- The CP-CP capable links exist between an end node and a network node, and the end node already has CP-CP sessions with a different network node.
- The CP-CP capable links exist between two adjacent end nodes.

If the CP-CP session is deactivated, you need to either reactivate the CP-CP session or deactivate all the CP-CP capable links and reactivate them with CPCP=NO.

In an APPN network you do not need meshed connectivity of CP-CP sessions with every node having a CP-CP session with every other node in the network. With APPN it is necessary only for each APPN node to have a CP-CP session into the network. A CP-CP session is not required between the end points of an LU-LU session. Defining CP-CP capable links to every other node in your network might cause excessive topology update flows. It is recommended that you define a primary and backup link with CONNTYPE=APPN and CPCP=YES for each node in your network. It is also recommended that you define additional APPN-capable links with CPCP=NO to provide efficient routing for your LU-LU sessions.

CP-CP sessions between two network nodes are used to perform searches for resources, exchange topology information, and can be used to register resources with a central directory server.

After an APPN connection has been established, identification information is exchanged between the nodes, and CP-CP sessions can be started between the control points in the directly attached nodes.

After the CP-CP sessions are established, the two nodes exchange CP capabilities, which indicate the level of network services provided by the control point. When both nodes are network nodes, they exchange topology database update (TDU) messages. The TDU messages contain identifying information, node and link characteristics, and resource sequence numbers that identify the most recent updates for each of the resources described in the TDUs.

When an LU-LU session is requested, information is transported across the CP-CP sessions to request that an LU-LU session be established between the two session partners. The control points that own the two session partners do not have to have a direct CP-CP session between them in order to establish the LU-LU session.

User sessions

User sessions are called LU-LU sessions in both subarea and APPN networks. An LU-LU session between two logical units enables one user to communicate with another. An example of an LU-LU session is an application, such as CICS, communicating with a workstation. LU-LU sessions are used only for user data, not for network control data.

A VTAM node can both initiate sessions and respond to session initiation requests. An LU-LU session is established when a BIND is sent from the primary LU to the secondary LU. Dependent LUs act as secondary LUs only and have an LU-LU session limit of one. Independent LUs can have parallel sessions between the same pair of LUs and can have multiple sessions between one LU and several other LUs.

Independent LUs exchange only UNBIND and RSP(UNBIND) for session termination. Independent LUs can be either the PLU or the SLU for any given session. An UNBIND can be sent by either LU.

A VTAM node can have both an SSCP and a CP. Therefore, VTAM nodes support the following types of session initiation over SSCP-SSCP or CP-CP sessions:

- Sessions initiated by the primary and secondary LU
- Queueing of session initiation requests
- Sessions initiated by a third party (CLSDST PASS)
- Orderly termination of existing sessions
- Automatic logon sessions
- Forwarding of session release requests

In an APPN network, the owning CP for an LU uses CP-CP sessions to obtain location and routing information needed to establish an LU-LU session. The owning CPs of the session partners do not have to be adjacent to each other and do not require having a direct CP-CP session with each other if they are not adjacent.

Sessions are established using a specific mode. The mode identifies session parameters and the APPN Class of Service (COS) name to use for the session. In subarea networks, the Class of Service indicates the virtual route and transmission priority to be used for this session. In APPN networks, the APPN Class of Service tables specify weights to be applied to the various routes in the APPN network based on the characteristics of the links and nodes. These weights are calculated, taking into account the various characteristics such as acceptable level of security, cost per byte, cost per connect-time, propagation delay, and effective capacity. For more information about Class of Service in APPN, see [“APPN class of service” on page 15](#).

VTAM allows dependent LUs, such as workstations and printers supporting LU types 0, 1, 2, and 3 in one VTAM domain to access VTAM application programs in a different domain by using APPN protocols between domains. All dependent LU types are supported through APPN networks with one exception: Sessions with binary synchronous 3270 devices are not supported over APPN links.

Examples of these dependent LU types are provided in [Table 1 on page 12](#).

<i>Table 1. LU types</i>	
LU type	Example
Type 0	3650 and 4700 financial terminals
Type 1	3767 and 3770 remote job entry (RJE) stations
Type 2	3270 interactive displays
Type 3	3270 printers

How VTAM locates resources

For VTAM to locate a resource in the network, VTAM must either have a definition for the resource or be able to dynamically define the resource. Following are descriptions of how resources are located in a subarea or APPN network.

Locating resources in a subarea network

In a subarea network, when VTAM receives a session initiation or directory search request for a resource that is not located in that VTAM domain, it attempts to locate the resource by sending the request to its adjacent SSCPs.

If the resource is statically defined as a CDRSC (or matches an active model CDRSC) and has the CDRM operand defined, VTAM routes the request to the SSCP specified in the CDRM operand. VTAM uses an adjacent SSCP table if:

- The CDRM operand is not defined
- The resource is not found at the CDRM specified
- VTAM has no SSCP-SSCP session with the specified CDRM
- The CDRSC is dynamically defined

The adjacent SSCP table can be statically defined or, if the start option DYNASSCP is defined as YES (which is the default), it can be dynamic. VTAM searches each SSCP in the table for the resource until the resource is found or the table is exhausted. See [“Adjacent SSCPs” on page 431](#) for additional information about using adjacent SSCP tables.

For independent LUs located on or through adjacent LEN nodes, VTAM uses these definitions to determine connectivity to the resource when establishing a session with the independent LU.

Locating resources in an APPN network

In an APPN network, when VTAM receives a session initiation or directory search request for a resource, VTAM checks its directory services database for the location of the resource. The directory services database contains information about the location of resources obtained through resource definition, resource registration, and as the result of successful searches. If the resource is not located in that VTAM database, VTAM issues a broadcast search to all other APPN nodes or a search to a central directory server, if one is available. If there are interchange nodes in the network, at the end of a broadcast search, they are sequentially allowed to search the subareas to which they are attached. VTAM uses a cross-subnetwork search to search across a subnetwork boundary.

Route selection

VTAM selects routes differently in subarea and APPN networks.

Routing in a subarea network

Defining a route in a subarea network involves the following elements:

- Subarea links
- Transmission groups (TGs)
- Explicit routes (ERs)
- Virtual routes (VRs)
- Path definitions
- Logon mode tables
- Subarea class of service (COS)

Subarea links are the different types of connections that can be defined between subareas. For example, links include a channel connection between two VTAMs.

Transmission groups (TGs) allow links with similar characteristics to be logically associated. Several links between the same two subarea nodes might have the same TG number, or the TG number might represent only one link.

Explicit routes (ERs) consist of a list of TGs and subarea nodes in one direction between the two end points of the route. For each ER in one direction, there must be a corresponding ER in the opposite direction that follows the same path.

Virtual routes (VRs) represent the complete route path. Each VR has a specific number and the VR definition associates that VR number with the ER to be used for a session. The VR specifies two one-way ERs to create a round-trip path.

Explicit routes and VRs are defined within path definition statements. Each path definition statement gives the ultimate destination point of the path. The ER gives the transmission group and subarea number for the next hop on the path. The association of the VR number with the ER within the path definition connects the individual hops together into a complete path from one end point to the other.

Each LU-LU session has an associated logon mode table entry. This entry specifies, among other things, the name of a subarea COS that contains a list of pairs of VR numbers and transmission priorities that can be used to establish the session. Using the subarea COS, you can specify the VR that best fits the characteristics needed for that session such as capacity, security, and speed.

For more information about subarea routing, see [“Network routing for subarea nodes”](#) on page 256.

Routing in an APPN network

APPN nodes provide a set of topology and routing functions that are used with the directory functions to find resources and calculate routes for sessions. Links must be defined to the nodes that own them, but routes over other links in the network can be learned dynamically.

Topology and routing services

The following concepts describe the topology and routing functions:

- Topology database
 - Local
 - Network
- Class of Service
- Route selection

Topology database

To enable a network node to provide routing functions to and from itself and the end nodes it serves, every network node maintains a network topology database that has complete and current topology information about the network. This topology information consists of the characteristics of all network nodes in the network and of all transmission groups (TGs) between network nodes. The end nodes in the network and the TGs connected to them are not considered network topology information.

This section describes the two kinds of topology databases in an APPN network and explains how each is used. Local topology information and network topology information is maintained at each network node. The local topology information is unique to the node; the network topology information is replicated at all network nodes.

Local topology database

A local topology database contains information about a node and about the TGs connecting the node to other type 2.1 nodes. Both end nodes and network nodes have local topology databases.

In an end node, the local topology database is used to calculate a route to the network node server and to supply the endpoint TG vectors (vectors that describe the TGs connected to the end node) to the network node server. In a network node, the local topology database includes information about the attached end nodes. The local topology database is not saved when VTAM is stopped and restarted.

Network topology database

The network topology database is referred to as the *topology database*. The network topology database is kept only in network nodes and contains complete and current information about all network nodes in the network, and about all transmission groups connecting them. It does not contain any information about LUs, end nodes, or LEN nodes.

The network topology database dynamically records changes in network topology and can be checkpointed to a file and reloaded when VTAM is restarted to reduce the amount of network traffic when a node is restarted.

To keep the topology database current, network nodes in an APPN network send each other topology database updates (TDUs) over CP-CP sessions whenever a resource (node or link) is activated, becomes inoperative, deactivated, or its characteristics change. Only the current changes are included in the TDU, not the complete topology. The TDUs contain information about the network nodes and

the TGs between the network nodes. Every network node receives the TDUs, so all have the same view of the network.

A resource sequence number (RSN) is associated with each resource, and is incremented each time that resource changes status. The RSN is used to determine when all network nodes have received the most current information about the node or TG, at which time the TDU activity for that resource ceases. As long as all network nodes follow the same rules with regard to TDU processing, the RSN guarantees that a TDU war does not occur. A TDU war is an endless exchange of TDUs in contention over the same topology resource, resulting in continuous performance degradation of the APPN network. See the Display TDU information section in [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#) for information about diagnosing TDU wars.

In an APPN network, several mechanisms are used to ensure that unnecessary or excessive TDUs are not propagated throughout the network. For example, multiple TDUs can be grouped and sent out together if there are updates to multiple local TGs.

The topology database is used by the network node to select routes for sessions that originate at the LUs in it and at the ENs that it serves.

APPN class of service

An APPN class of service (COS) defines the required or requested characteristics of a route for a session. A COS consists of a set of ranges of acceptable values for the characteristics of links and nodes to be used for a session specifying that particular COS. APPN Classes of Service are defined in a VTAMLST definition list.

Unlike the COS for the subarea network, where the COS is actually a list of VRs that are acceptable for a particular COS, APPN COS specifies the types of routes that are acceptable for a Class of Service.

Class of Service database

The COS database exists on all VTAM network nodes (NNs) and end nodes (ENs). It is the same on both ENs and NNs, and it is built when VTAM is initialized.

The COS database contains the following items:

- COS definitions from VTAMLST definition files

IBM provides three sets of COS definitions: COSAPPN, ISTACST2, and ISTACST3. Each set contains the same seven default APPN COS definitions. However, differences exist in the way the seven Classes of Service are defined in each set. See [“What are the IBM-supplied default classes of service?” on page 243](#) for information about the differences between the three sets and how to determine which one is appropriate for your configuration. It is recommended that you use the defaults wherever possible and that the COS definitions be the same throughout the network. The default APPN COS definitions in z/OS Communications Server (COSAPPN) are the same default Classes of Service provided by other APPN products.

One or all of the IBM-supplied sets of COS definitions can be included in VTAMLST. User-defined Classes of Service can also be included. For information about how to define your own COS, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

You can have only one set of COS definitions in VTAMLST active at any time. The IBM-supplied COS definitions in APPNCOS are automatically activated during z/OS Communication Server initialization. COS definitions can also be activated with a VARY command or by specifying the name of the VTAMLST member in a CONFIG list. Entries can be added or modified by activating or reactivating the VTAMLST definitions.

- A list of logon mode names and the corresponding APPN COS

A set of default mappings for the seven IBM-supplied classes of service is included in the default logon mode table. It is recommended that users also code one of the seven standard APPN COS names in any user defined mode tables. For information about defining user-specified modes and corresponding Classes of Service, see [z/OS Communications Server: SNA Resource Definition Reference](#).

- A list of calculated routes

As an NN calculates routes for particular classes of service, the routes are stored in VTAM memory. Storing routes allows them to be reused, rather than being recalculated each time a route is requested.

Routes are periodically recalculated to allow for distribution of sessions over equally weighted paths. Start options control this process; see [z/OS Communications Server: SNA Resource Definition Reference](#).

Route calculation and selection

VTAM provides automatic, preferred route selection based upon dynamic network characteristics and requested COS. During session establishment, the network node server of the origin LU refers to the topology database to calculate and select the current best route through the APPN network from the primary LU to the secondary LU for the requested COS.

An NN calculates routes for sessions that originate at the LUs in it and at the ENs it serves. When a route is calculated, it is stored and can be reused. The stored routes are unique to the node and are never broadcast. The cache of stored routes is maintained automatically. When the maximum cache size (set by the NUMTREES start option) is exceeded, the least recently used routes are discarded first.

Route selection is based on how the actual characteristics of each node and TG along the possible paths match the characteristics required by the requested COS. The route that an NN selects is the current least-weight (or best) route from the node containing the origin LU to the node containing the destination LU.

Addressing

VTAM uses different addressing schemes in subarea and APPN networks.

Subarea addressing

Subarea addressing uses an addressing scheme that consists of a subarea number representing the subarea a resource belongs to, and an element address, which uniquely identifies that resource within the subarea. The subarea number for VTAM is given by the HOSTSA start option. (The subarea number for NCP is determined by the SUBAREA operand on the BUILD statement.) The element address is assigned from a pool of addresses. Both VTAM and NCP maintain a pool of addresses for their attached resources.

APPN addressing

APPN nodes use local form session identifiers (LFSIDs) instead of subarea addresses for external communication in an APPN network. VTAM APPN nodes have subarea addresses that are completely internal to the nodes, used by the NCP in a composite network node to communicate with its owning VTAM. This subarea address is never used outside the composite network node.

When High-Performance Routing (HPR) is in use, LFSIDs are not used for the HPR portion of the sessions. Instead, automatic network routing (ANR) labels are used by each node to route the HPR network layer packets (NLPs). For further information about HPR, see [“High-Performance Routing \(HPR\)” on page 388](#).

Controlling network data flow using pacing

Pacing is a means of controlling the flow of messages in the network to avoid congestion. Data congestion results whenever the rate at which data is going into a network exceeds the capacity of the network. Response times might lengthen and throughput might decrease. Severe or prolonged congestion in one part of a network can affect the other parts, causing overall network efficiency to suffer.

VTAM monitors traffic and limits congestion in the network with pacing. Pacing controls network flow by limiting the amount of data a transmitter can send before receiving an acknowledgment from the receiver. For example, if the agreed amount of data is three units, after sending three units of data the transmitter must wait for the receiver to send an acknowledgment before sending any more data.

Table 2 on page 17 lists the two levels of pacing that are involved in the flow-control process.

Table 2. Pacing types	
Pacing type	Description
Session	Session pacing prevents overrun of the buffers of the LUs at the session endpoints. There are two types of session pacing: fixed and adaptive. See “Virtual route pacing” on page 270 for more information.
Virtual Route	Virtual route pacing is used by subarea nodes to avoid buffer congestion at the virtual route endpoints. See “Virtual route pacing” on page 270 for more information.

Both pacing mechanisms allow pacing to occur inbound and outbound independently. Pacing values do not have to be the same for the two directions. The term *inbound pacing* is associated with the pacing window for message units being received by a node. *Outbound pacing* is associated with the pacing window for messages being sent by the node.

Pacing concepts

Before continuing with the specifics of the two pacing mechanisms, you should understand basic pacing concepts. Pacing is controlled using pacing window techniques. A *pacing window* is the number of message units that a pacing endpoint can send to another network component at one time. The first message unit contains a pacing request, which indicates that a pacing response must be sent by the receiver before another pacing window of message units can be sent.

Pacing window

For example, in [Figure 1 on page 17](#) the pacing window size is three. When Node A sends three message units, it must wait for a pacing response (from Node B) before sending another message unit. The pacing response restores the pacing window size.




Figure 1. Pacing flow – outbound pacing

Figure 1 on page 17 shows pacing from the sender's outbound pacing perspective. There could also be a pacing window size associated with the sender's receiving message units (inbound pacing). The pacing window size can be different from the outbound pacing window size.

In [Figure 2 on page 18](#), Node A inbound pacing window size is two. Node B, after sending two message units, will not send additional message units until Node A sends a pacing response to Node B as shown in [Figure 2 on page 18](#).




Figure 2. Pacing flow – inbound pacing

A pacing window response can be sent before the pacing window size is reached. If the receiver sends the pacing response early (before having received the maximum number of requests), the sender transmits the first group of requests and can then send another group of requests.

Figure 3 on page 18 shows the flow when a pacing response is sent before reaching the pacing window size. Node A inbound pacing window size is three. Node A could send the pacing response after receiving message unit 1. Node B could send up to three more message units before another pacing response is required to be sent by Node A.




Figure 3. Pacing flow – receiving early pacing response

Chapter 3. Implementing a VTAM network

This topic explains how to implement a VTAM network and includes the following topics:

- Using start options and configuration lists
- Identifying resources to VTAM
- Verifying a VTAM network

Implementing a VTAM network involves the following tasks:

- Use of start options and an optional configuration list
- Identification of particular resources in the network to VTAM
- Identification of any paths required for network routing
- Establishment of sessions
- Operation of VTAM itself

You will also need to understand the VTAM start procedure and the associated data sets. For information about this process, see [z/OS Communications Server: New Function Summary](#).

This topic gives descriptions of VTAM network requirements and an example of implementing a VTAM network and includes the following subtopics:

Start options

Start options control the conditions under which VTAM runs. No matter what network configuration you are implementing, you need to code some start options. Start options are coded in files named ATCSTRxx, where xx specifies the identifier of a particular start option file.

Configuration list

A configuration list lets you specify which resources are activated automatically when you start VTAM. You are not required to define a configuration list, but it makes the VTAM operator's job easier because VTAM activates the resources in the configuration list automatically. Configuration lists are coded in files named ATCCONxx, where xx specifies the identifier of a particular configuration list.

Application programs

Each host application program you are running must be defined to VTAM as an application program minor node within an application program major node. A major node, such as an application program major node, is a set of minor nodes (in this case, the application programs) that can be activated and deactivated as a group. For more information about application programs, see [Chapter 13, "Application programs,"](#) on page 277.

Note: An application program may become a shadow resource if a CDRSC with the same name already exists when the major node containing the application program's definition is activated. For more information about shadow resources, see ["Shadow resources"](#) on page 438.

Subarea nodes

If you have one or more NCPs in your network, define them in NCP major nodes and channel-attachment major nodes.

If there are other VTAMs in your network, define connections to them in NCP major nodes, channel-attachment major nodes, or external communication adapter (XCA) major nodes. If your VTAM is to have SSCP-SSCP sessions with other VTAMs, also create a cross-domain resource manager major node and minor nodes for your VTAM and adjacent VTAMs with which your VTAM is to have such sessions.

You need to define paths for data flow between VTAM and any owned NCPs, and for data flow to and from other VTAMs and NCPs in your network over subarea connections. Paths are defined in PATH definition statements.

APPN nodes

If VTAM is going to use APPN functions, specify the NODETYPE start option.

Peripheral nodes

Other physical devices in your network must also be defined to VTAM, whether they are directly attached to the host or to an NCP. Peripheral nodes are dynamically defined or manually defined in channel-attachment, external communication adapter, local non-SNA, local SNA, LU group, model, packet, and switched major nodes. Logical units in or attached to peripheral nodes are defined with LU statements along with the peripheral nodes in the major node definitions, or if they are independent LUs, they can be defined in cross-domain resource major nodes. The major node you choose depends on the characteristics of the device you are defining.

Figure 4 on page 20 is a sample VTAM subarea network. The host is running MVS, VTAM, and two VTAM application programs. NCP11 is a channel-attached 3745 controller, and NCP12 is a link-attached 3725 controller. Peripheral nodes are connected using leased lines.




Figure 4. VTAM network

You need to code at least the following resources to define the network in Figure 4 on page 20:

- Start options describing the host (NETID, SSCPID, and SSCPNAME).
- Application program minor nodes defining APPL1A and APPL1B.
- Two NCP major nodes defining NCP11 and NCP12. These include the NCP definitions. In the NCP major node, also code GROUP, LINE, PU, and LU definition statements to define the peripheral nodes attached to the NCPs.

Note: APPN nodes, PUs, and LUs can be defined dynamically.

- PATH definition statements in each subarea to attach to the other subareas. For example, in NCP11, you would code two PATH definition statements to attach to subarea 1 (DESTSA=1) and subarea 12 (DESTSA=12).

Using start options and configuration lists

When you start VTAM, you can enter start options on the START command, or you can write one or more start lists to assist you in starting VTAM. You can also activate your resources either manually or automatically. To automatically activate some or all of your resources, you can write one or more configuration lists. Writing start lists and configuration lists:

- Reduces the amount of operator involvement and the chance of entering incorrect information
- Enables VTAM to initialize the domain faster

You can code multiple start options or configuration lists and use these lists according to your needs. After you specify start options and configuration lists, they remain in effect until you modify them or until VTAM is halted.

The syntax rules for start lists and configuration lists are less restrictive than those for other VTAM definition files. Full line comments must always begin with an asterisk (*) in column 1 and cannot be continued on the next line using a continuation character in column 72. No continuation characters are required in column 72, and commas are not required between statements. Continuation characters and commas are permitted, if required, and will not cause a problem.

For more information about coding start options and configuration lists, see the [Rules for coding start option lists and configuration lists](#) section of the [z/OS Communications Server: SNA Resource Definition Reference](#).

Start options

This section contains information about the following topics:

- Required start options
- Recommended start options
- Sources of start options
- Start option processing
- Creating start option lists
- Improving VTAM performance using start options
- Increasing Host Subarea Element Addresses
- Enhanced addressing for session managers
- Diagnosing problems using start options

Start options provide information about the conditions under which VTAM runs. They also enable you to tailor VTAM to meet your needs each time VTAM is started. Many operands can have defaults specified as start options, thus reducing the amount of coding required. Many start options can be dynamically modified and also displayed. A complete list of start options is listed in the [z/OS Communications Server: SNA Resource Definition Reference](#).

Required start options

The following three start options are required:

SSCPID

The SSCPID start option provides VTAM with a unique numeric identifier. The SSCPID value is used by some physical units to identify the VTAM with which it is in session.

If you plan to expand or incorporate a single-domain network into a larger network, be sure that the value of SSCPID is unique for each host. The SSCPID value you specify must be different from the SSCPIDs in other networks that can be in session with this host.

If you plan to incorporate a non-VTAM host, the values of the VTAM SSCPID and HOSTSA start options should be coded so that one of the following items is true:

- The values for both SSCPID and HOSTSA are higher than the non-VTAM host.
- The values for both SSCPID and HOSTSA are lower than the non-VTAM host.

When activating the cross-domain resource manager-cross-domain resource manager (CDRM-CDRM) session, the host with the higher SSCPID value sends the activate cross-domain resource manager (ACTCDRM) and the host with the higher HOSTSA value activates the route. Unpredictable results occur if the route chosen for the ACTCDRM is different from the route activated.

SSCPNAME

The SSCPNAME start option provides a unique name for VTAM. This option is required for a single-domain network, but is primarily used in multiple-domain and multiple-network environments to identify a particular VTAM. The SSCPNAME option must be different from the HOSTPU start option that identifies the physical unit within VTAM.

The SSCPNAME option is also used as the CP name for a VTAM that implements APPN. Although the SSCPNAME and CP name are identical, note that the SSCP and CP are distinct resources and can be displayed individually.

Note: The SSCPNAME should match the name that is coded in the CDRM major node for this VTAM.

NETID

The NETID start option provides VTAM with the network identifier. If you connect your VTAM to another network, the network identifiers must be unique. The network name should conform to an overall naming standard, such as:

- The first two characters identify the country in which the network is managed and comply with the International Standards Organization (ISO) Standard 3166.
- The next four characters identify your enterprise.
- The last two characters identify a particular network within your enterprise.

For example, an IBM network within the United States could have the network identifier *USIBMCMK*, where *US* identifies the country, *IBMC* identifies the enterprise, and *MK* identifies a particular network within IBM.

Notes:

1. Using the # symbol for defining network identifiers (NETIDs) is not recommended.
2. To assist SNA network owners in controlling the uniqueness of their network IDs, IBM has established a worldwide SNA Network Registry. Registering a network ID in the SNA Network Registry confirms its uniqueness from all other network IDs within the registry, thus minimizing future connectivity problems. To register a network ID, contact your IBM branch office representative.

Recommended start option

The HOSTPU start option is recommended (but not required) for identifying VTAM to the network.

Use the HOSTPU start option to assign a user-defined name to the VTAM host physical unit. The default name for the host physical unit is ISTDPU. Because the NetView program uses this name to associate network information to a specific VTAM, you should make sure that this name is meaningful and unique throughout the network.

For start options you are not required to code, VTAM uses the default values. You can override any start option value when VTAM is started or enter new start option values at the operator console.

REUSASID=YES is an optional but recommended start option to add at the end of the start option parameter list. This option helps to maximize storage usage. REUSASID specifies that the address space attempts to use an ASID from the reusable address space pool.

Sources of start options

When VTAM is started, you can provide the options from any combination of the sources listed below (arranged from the lowest to highest priority):

- IBM-supplied values internal to VTAM (default values).
- Default start option list. When you start VTAM, a list of user-defined default values (ATCSTR00) is read from the VTAM definition library. If VTAM cannot find the ATCSTR00 file, it prompts you, giving you a choice of:
 - Using default values

- Using an alternate start option file (ATCSTRxx, where xx specifies the identifier of the start option file you want to use)
- Canceling the start attempt

If you do not want to use the ATCSTR00 file, you can avoid the prompt by coding an ATCSTR00 file that contains only comments. VTAM issues a message that the file is empty, but does not interrupt start processing.

- Supplemental start option list. The VTAM operator can enter the LIST=xx start option on the START command to name another list to supplement ATCSTR00. The supplemental list (ATCSTRxx) can override options in the default ATCSTR00 list, and the IBM-supplied internal default values.
- Backup start option list. The LISTBKUP start option can be placed in a start list (ATCSTR00 or ATCSTRxx) to indicate the processing that occurs if a start option in the list is not valid. You should place the LISTBKUP start option first in a start list, so that it will have been processed if an error is found for another start option.
- Start options entered by the operator. The VTAM operator can enter additional start options on the START command, and also during VTAM startup if prompted. (VTAM prompts for start options if the operator does not enter any start options on the START command.) To prevent VTAM from prompting for start options during startup, code the NOPROMPT option in ATCSTR00.

Start options entered by the operator can override options in a specified supplemental start option list, and start options in the default ATCSTR00 list and the IBM-supplied internal default values.

To enter a list of options longer than the length of the console, place a comma after the last start option that will fit on the console and VTAM will continue to prompt for more start options.

- Start options reentered by the system operator to correct errors in previously processed start options.

Start option processing

In “Sources of start options” on page 22 the list of start option sources is arranged from the lowest priority to the highest. VTAM processes the start options with the lowest priority first and those with the highest priority last. If conflicting specifications exist for a particular start option, the last valid specification entered always overrides any previous specification. For example, if ATCSTR00 contains CONFIG=01, but the operator enters LIST=ST and ATCSTRST contains the start option CONFIG=02, VTAM activates the configuration defined by ATCCON02.

If VTAM detects an error when reading in a start option list, it notifies the operator. The operator must then do one of the following actions:

- Respond to the error message by selecting the "continue processing" option and then reenter any start options that were specified incorrectly in the start option list. The operator can also enter additional options at this time.
- Specify a different start option list to use by entering xx (for the required ATCSTRxx start option list). This cancels the start attempt and then begins start processing again with another start option list.
- Halt VTAM. This cancels the start attempt so that you can correct the start option list and try again.

After reading in a start option list, if VTAM detects an error while processing a specific start option, the following condition occurs:

- Backup processing occurs if the LISTBKUP start option was processed and the start list in error is not a backup start list (only one backup start list can be processed).
- If the LISTBKUP start option was not processed, the operator is provided with the same three choices available when VTAM encounters an error while reading in a start list.

Use the LISTBKUP start option to specify:

- Another start list should be used in place of the start list in error (LISTBKUP=xx).
- All start option values existing before the file in error was processed should be used (LISTBKUP=DEFAULTS).

- All valid start options in the file in error will be processed, and incorrectly specified start options should remain as previously set until VTAM reprompts for these options (LISTBKUP=PROMPT).

Creating start option lists

To use a start option list, create a list and put it in a partitioned data set member named ATCSTRyy in SYS1.VTAMLST. The yy value must be the same as the yy value in the LIST=yy start option entered by the operator.

When the VTAM operator enters LIST=yy as a start option, VTAM first attempts to locate ATCSTR00. If ATCSTR00 does not exist, VTAM sends a warning message to the operator. To avoid receiving this message, create an ATCSTR00 file that contains only comments or start options that are always used for that particular VTAM.

Following is a sample start list ATCSTRyy.

```
SSCPID=01
SSCPNAME=SSCP1A
NETID=NETA
DYNASSCP=YES
SSCPDYN=YES
HOSTSA=1
```

The operator can start VTAM with this list using LIST=yy on the start command.

One start list cannot see another using LIST=yy. Only the operator can enter LIST=yy to use a start list.

Improving VTAM performance using start options

You can use the following information to improve performance using start options:

- VTAM search reduction support can be used with the SRCHRED, SRTIMER, and SRCOUNT start options.

Your use of start options also affects VTAM use of host processor time and storage. Buffer pool specifications, the limit of buffer usage for session outage notification, and the number of subtasks used to perform dump, load, and restart processes impacts VTAM host resources. For more information about these start options and other factors that impact VTAM performance in your network, see [Chapter 20, “Tuning VTAM for your environment,”](#) on page 499.

Reducing network search overhead using search reduction

When a resource is unreachable in a network, futile attempts to reach it can still occur. Excessive searching for unreachable resources can adversely affect network performance. Therefore, VTAM provides search reduction support, which limits requests for unreachable resources.

If search reduction support is enabled, an APPN or subarea search reduction entry is created when a resource discovery search (RDS) fails to locate a target resource. An APPN search reduction entry is saved as a directory entry in the directory services database. A subarea search reduction entry is saved as a CDRSC in the subarea resource definition table (RDT). VTAM then reduces searching for the unreachable resource in the appropriate part of the network during a designated amount of time or until a designated number of requests have been received for the resource.

If a subarea search reduction entry exists for a destination resource and VTAM is the SSCP of the originating LU (including instances where VTAM is an interchange node for requests entering subarea from APPN), search reduction optimizes adjacent SSCP table searching for that resource. VTAM limits attempts during the designated period to locate the unreachable resource in the subarea network.

If an APPN search reduction entry exists for a destination resource, the scope of optimization varies depending on the node type of VTAM.

- If VTAM is the central directory server or network node server for the originating LU, including instances where VTAM is an interchange node for requests entering APPN from a subarea, search reduction optimizes APPN searching in VTAM's subnetwork. VTAM limits attempts during the designated period to

locate the unreachable resource in its APPN subnetwork. For information about of APPN subnetworks, see [“APPN multiple network connectivity”](#) on page 75.

- If VTAM is a border node for requests received across a subnetwork boundary, search reduction optimizes APPN searching in the VTAM subnetwork. VTAM limits attempts during the designated period to locate the unreachable resource in its APPN subnetwork.
- If VTAM is a border node for the originating LU, search reduction optimizes APPN searching throughout the APPN network. VTAM limits attempts during the designated period to locate the unreachable resource in the APPN network.

To use VTAM search reduction support, code the SRCHRED=ON start option or enter the MODIFY VTAMOPTS command with the SRCHRED=ON operand. Use the SRTIMER start option, or MODIFY VTAMOPTS with the SRTIMER operand, to designate the amount of time that VTAM should limit requests for a resource. Use the SRCOUNT start option, or MODIFY VTAMOPTS with the SRCOUNT operand, to designate the number of requests for a resource that VTAM should limit before attempting to locate the resource again.

In determining a proper value for SRTIMER or SRCOUNT, keep in mind that the value should be large enough to limit searching for an appropriate amount of time, but small enough so that a resource that becomes available again does not remain unreachable because of waiting for SRTIMER or SRCOUNT to expire. Also, keep in mind that the SRCLEAR=YES operand can be used on the MODIFY RESOURCE command to clear search reduction entries. That is, you can use this when a resource becomes available to clear search reduction entries, so that the specified resource can now be located.

Note that SRCLEAR=YES only clears search reduction entries; it does not turn off search reduction support. If a subsequent RDS fails to locate the target resource, VTAM creates a new search reduction entry and limits subsequent requests for the unreachable resource during the designated amount of time or until the designated number of requests have been made.

Notes:

1. The SRTIMER and SRCOUNT operands can also be specified for resources on the CDRSC definition statement or the GROUP definition statement in a CDRSC major node. These values can be dynamically modified using the MODIFY RESOURCE command.
2. The DISPLAY DIRECTORY command displays information about search reduction entries in the directory services database. The DISPLAY ID command displays information about search reduction entries for a particular CDRSC, and entries in the directory services database (if the IDTYPE=RESOURCE operand is used).

When a resource has search reduction information associated with it, these commands display the defined values for SRTIMER and SRCOUNT, and the remaining amount of time and remaining number of searches before VTAM discards the search reduction information.

3. Use DISPLAY VTAMOPTS for information about the settings of the search reduction start options.
4. When using SRTIMER, SRCOUNT, or both, along with AUTOTI for autologon sessions, an unsuccessful autologon-session search for a controlling PLU can cause the creation of a search reduction entry for it. In this situation, be aware of the following situations:
 - If AUTOTI=x and SRTIMER=y, where y is greater than x, when the AUTOTI timer expires, subsequent requests for the PLU are limited by the search reduction entry until SRTIMER also expires.
 - If AUTOTI=x and SRCOUNT=z, when the AUTOTI timer expires, subsequent requests for the PLU are limited by the search reduction entry for (x * z) seconds. At that time, the search reduction entry expires.

See [“Controlling searches”](#) on page 413 for additional information.

Increasing host subarea element addresses

For APPN resources, VTAM maintains a table of network addresses used only by VTAM. To increase the number of APPN sessions VTAM can support, VTAM allows extended host element addressing using the ENHADDR=YES start option. If ENHADDR=YES is specified, high-order addresses are assigned for the primary LU (and often the secondary LU as well) when the addresses of both the primary LU and

secondary LU are assigned by the same VTAM (not by an NCP). This includes only sessions where both session partners are either a VTAM application or an LU that is attached to VTAM (not to an NCP). The maximum number of extended address elements that can be assigned is 33 161 216. To monitor the number of extended addresses currently in use at any time, use the `D NET,STATS,TYPE=VTAM` command. Stat 165 displays such a count in decimal format.

To further extend the usage of enhanced host element addressing, high-order element addresses are assigned for EE PUs, EE lines, RTP PUs, and DLUR-served PUs (when possible). This is done regardless of the coding of the ENHADDR start option.

Note: The extended and base address in use counts (stats 164 and 165 from the `D NET,STATS,TYPE=VTAM` command) might be one greater than the respective highwater values (stats 161 and 162 from the `D NET,STATS,TYPE=VTAM` command). This is because the addresses specified by the highwater values are actual addresses from a pool beginning with value 0, while the address counts are actual counts of addresses in use from that pool of addresses.

Enhanced addressing for session managers

To further extend enhanced host element addressing, applications that are activated within an application program major node can be assigned a high-order base address. When an LU-LU session begins for the application, the addressing capabilities of the partner LUs are assessed. If either partner does not permit use of a high-order address, a low-order auxiliary address is used.

The extent to which high-order addresses are used is determined by the parallel session capability of the application and the type of host that the application belongs to.

If the host is not pure APPN (in other words, supports subarea), applications that are not parallel session-capable will be assigned a high-order address when the application major node is activated and a low-order address is acquired when an LU-LU session involving the application requires a low-order address.

Note: Session manager applications are not typically parallel session capable.

If the host is pure APPN, all application addresses will be high-order addresses regardless of its parallel session capability.

This enhancement increases the number of addresses available for sessions and reduces the instance of sessions failing because of insufficient network addresses.

Diagnosing problems using start options

The following start options provide help with problem diagnosis. You can implement these when you start VTAM.

ASIRFMSG

Enables you to specify the SSCPs in which to issue messages IST890I and IST896I when an autologon session initiation request fails.

BSCMDRS

Enables you to specify which BSC miscellaneous data records are to be recorded in LOGREC.

DSIRFMSG

Enables you to specify the SSCPs in which to issue messages IST663I, IST664I, IST889I, and their message groups when searches that might not result in a session fail to locate the target resource.

Recommendation: Because of the volume of messages that can be generated, it is recommended that you disable this option during normal operation. It is recommended that you enable this option (using the `MODIFY VTAMOPTS` command) on all necessary hosts only when trying to diagnose specific problems. After the problem has been diagnosed or documentation has been collected, you should again disable this option (using the `MODIFY VTAMOPTS` command).

Tip: When SIRFMSG is coded on CDRSC definition statements, the DSIRFMSG function (for DLU resources) and SIRFMSG function (for OLU and DLU resources) are run. This can be helpful in diagnosing problems involving those CDRSC resources.

ESIRFMSG

Enables you to specify the SSCPs in which to issue messages IST890I, IST892I, and IST893I when a session initiation request fails and extended sense data exists.

FSIRFMSG

Enables you to specify the SSCPs in which to issue additional search routing failure messages when sense data exists. When used in conjunction with the SIRFMSG or DSIRFMSG option, the FSIRFMSG option controls whether messages IST894I, IST895I, IST1704I, and IST1705I are issued. When used in conjunction with the LSIRFMSG option, the FSIRFMSG option controls whether messages IST2208I, IST2209I and IST1942I through IST1953I are issued.

INOPCODE

Enables a specific inoperative condition, if experienced by any resource, to take a VTAM dump.

Restriction: Both the resource (INOPDUMP) and condition (INOPCODE) must be enabled in order for a dump to be taken.

INOPDUMP

Enables a resource, if and when it experiences an inoperative condition, to take a VTAM dump.

Restriction: Both the resource (INOPDUMP) and condition (INOPCODE) must be enabled in order for a dump to be taken.

IOINT

Enables you to request that VTAM inform the operator when a VTAM request is not answered after a specified interval of time. After the period of time specified by IOINT, VTAM issues a message to the operator. If VTAM issues the message multiple times for the same request or several different requests, initiate the appropriate diagnostic actions to resolve the problem.

LSIRFMSG

Enables you to specify the network nodes in which to issue messages IST663I, IST664I, IST889I, and their message subgroups when APPN locates fail to locate the target resource.

Recommendation: Because of the volume of messages that can be generated, it is recommended that you disable this option during normal operation. It is recommended that you enable this option (using the MODIFY VTAMOPTS command) on all necessary hosts only when trying to diagnose specific problems. After the problem has been diagnosed or documentation has been collected, you should again disable this option (using the MODIFY VTAMOPTS command).

Tip: When SIRFMSG and CPNAME are coded on CDRSC definition statements, the SIRFMSG setting will override the LSIRFMSG start option (for DLU resources). This can be helpful in diagnosing problems involving those CDRSC resources.

MSGMOD

Enables you to see the name of the module from which a message is issued. If you specify MSGMOD=YES, VTAM inserts the last five characters of the module name that issued the message, following the VTAM message identifier (for example, ISTnnnI). However, using MSGMOD causes the last five characters of the message to be truncated if the message exceeds the maximum console message length allowed by the operating system.

NMVTLOG

Enables you to specify whether error records are written to the operating system's error data set (LOGREC). If NMVTLOG=NPDA (the default), VTAM records alerts in LOGREC only when the NetView hardware monitor, NPDA, is not active. If NMVTLOG=ALWAYS, VTAM always records alerts in LOGREC. If NMVTLOG=NEVER, VTAM never records them.

PPOLOG

Enables you to specify whether to send operator commands and their resulting messages to the primary program operator for logging.

PSSTRACE

Enables you to specify whether to record IRB and SRB entries in the VTAM internal trace table when the PSS trace option is in effect.

RSIRFMSG

Enables you to specify the SSCPs in which to issue messages IST1460I, IST1461I, IST2102I, IST2103I, and IST2104I to display the RSCV when a session initiation request fails and RSCV data exists.

SDLCMDRS

Enables you to specify whether SDLC statistical miscellaneous data records (MDRs) are written to LOGREC. If SDLCMDRS=YES (the default), VTAM records SDLC statistical MDRs in LOGREC.

SIRFMSG

Enables you to specify the SSCPs in which to issue messages IST663I, IST664I, IST889I, and their message groups when a session initiation request fails.

Note: SIRFMSG can also be coded on the APPL and CDRSC definition statements to help diagnose problems involving those resources. Coding SIRFMSG on a CDRSC definition statement also enables the DSIRFMSG function and the LSIRFMSG function (if CPNAME is also coded) for that resource.

SNAPREQ

Enables you to specify the number of requests for VTAM buffers between snapshot dumps.

TNSTAT

Enables you to specify whether, where, and how frequently to record tuning statistics.

TRACE

Enables you to collect information, such as data that flows through VTAM buffers, data transmitted on NCP links, and VTAM internal trace information. For more information about the trace facility, see [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#).

See [z/OS Communications Server: SNA Operation](#) for a list of the start options and descriptions of the operator commands that you can use to start VTAM trace facilities. The ASIRFMSG, ESIRFMSG, FSIRFMSG, and RSIRFMSG start options produce messages only when the SIRFMSG, DSIRFMSG, or LSIRFMSG start options are also in use. For descriptions of the start options, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Configuration lists

A configuration list specifies the resources that are to be activated when you start VTAM. Place the member names of the resources you want to have activated when VTAM starts into an ATCCONxx member in the VTAM definition library, where xx is any two alphanumeric characters. The value xx can then be used on the CONFIG operand of the VTAM START command, or on the CONFIG start option in your start option list, to specify which definitions are to be activated at startup. In [Figure 5 on page 28](#), to use the configuration list defined in ATCCON01, you can specify CONFIG=01 on the VTAM START command.




Figure 5. Configuration list

ATCCON01 contains the member names LOCSNA1, LOCNS1, NCP1, which will be activated when VTAM starts. CHAN1 and RSNA1 (found in the VTAM definition library) are not included in ATCCON01, and therefore must be activated manually by the VTAM operator using the VARY command, if they are to be used. If no configuration lists exist, the VTAM operator must use the VARY command to activate all resources.

Following is sample coding for configuration list ATCCON01.

```
LOCSNA1
LOCNS1
NCP1
```

Examples of resource definitions that can be included in a configuration list are:

- Paths
- Major nodes
- Minor nodes defined in previously listed major nodes
- Tables
- Dynamic reconfiguration files
- Network node server list

You can also use a configuration list to activate an NCP using a different PU name than the one specified in the NEWNAME operand on the BUILD definition statement. For example, if you want to load a communication controller with load module X3745A and want the PU name to be NCPA1, you would code the following in your configuration list:

```
NCPA1/X3745A
```

This would have the same result as using the following command:

```
VARY NET ACT, ID=NCPA1, LOADMOD=X3745A
```

Each of the names specified in the configuration list must be one to eight characters (one to seven characters for the load module name when specifying pu_name/load_module). The first character of each name must be alphabetical (A–Z) or one of the national characters (@, #, or \$). The remaining characters of each name must be alphabetical (A–Z), numerical (0–9), or one of the national characters (@, #, or \$). If a name is found that is not valid, VTAM stops processing the configuration list and is unable to initialize any of the subsequent major nodes.

Identifying resources to VTAM

In addition to specifying start options and coding configuration lists, you'll need to identify resources in the network to VTAM. Depending on your network, you might need to define a combination of the following resources:

- Application programs
- Adjacent control points
- NCPs
- Peripheral nodes
- LUs
- Other VTAMs

Any of these resources may be predefined to VTAM using a static definition statement. For example, applications are predefined using an application major node and application (APPL) statements; switched PUs are predefined using the NCP major node for the lines and switched major node for the PUs and LUs. In addition, many VTAM resources may be dynamically defined as VTAM learns of them. A switched PU may be defined when it dials in; an adjacent CP may be defined when the connection between the nodes is activated.

Coding concepts

This section describes basic coding concepts needed to define resources. For further information, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

To define your resources:

1. Code VTAM and NCP definition statements to define your major nodes.

The following restrictions apply to coding your definition statements:

- Names must be unique within a network.
- Names cannot begin with "IST" or any of the names "VTAM", or "VTAMSEG", or "VTAMSG2". Also, the name "VTAMTERM" cannot be used.
- Code each minor node and major node name in column 1.
- Code each definition statement in column 10.
- Code each operand in column 16 or one space after the definition statement.
- If your operands are longer than one line, code a continuation character in column 72 and continue operands in column 16 of the next line.

Following is a sample definition file KIMAP VTAMLST (the scale is included to show coding restrictions, such as the continuation character in column 72):

```
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...
KIMAP      VBUILD TYPE=APPL
AP1        APPL  AUTH=(PASS,ACQ) ,           X
           ACBNAME=ECHO
AP2        APPL  ACBNAME=ECHO
```

2. If you are defining NCP, code the NCP-only definition statements and operands as described in the NCP resource definition documents.
3. Generate the NCP as described in the *NCP, SSP, and EP Generation and Loading Guide*. If you are updating VTAM-only definition statements or operands, update the NCP definitions in VTAMLST; no NCP generation is required. If you are updating NCP-only definition statements or operands, update the NCP generation definition and generate NCP again, or use dynamic reconfiguration facilities.
Note: If you need to make changes before you can regenerate the NCP, VTAM and NTuneNCP provide functions to allow some NCP definitions to be dynamically changed.
4. Store your definitions in the proper VTAM library. You can concatenate VTAMLST files, but either the block sizes must be the same for all files or the file with the largest block size must be the first file in the concatenated list.

Sift-down effect

The *sift-down* effect enables you to code an operand on a higher-level node so that you do not need to recode it on each lower-level node for which you want the same value. As a result, the sift-down effect greatly simplifies the coding process.

The *z/OS Communications Server: SNA Resource Definition Reference* identifies and describes the definition statements and operands to which sifting applies. For information about definition statement sequencing and the sift-down level for each NCP operand, see the NCP resource definition documents.

Using MVS system symbols

Multiple virtual storage (MVS) system symbols enable you to reduce the number of VTAM definitions you must code in VTAMLST. These definitions include:

- Start option lists
- Configuration lists
- Major nodes and minor nodes
- Routing and dynamic reconfiguration definitions
- The following user-definable tables:
 - APPN class of service (COS) definitions
 - APPN-to-subarea COS mapping table

- Associated LU table
- Message-flooding prevention table
- Model name table
- Subarea-to-APPN COS mapping table
- SAMAP

Substitution text for the symbols is defined in MVS and is substituted by MVS in place of the symbols during VTAM startup, major node activation, and table activation.

By using MVS system symbols in VTAMLST, you can code a single start option list and a single configuration list that can be used to start VTAM on multiple systems. You can also code one set of major nodes, routing and dynamic reconfiguration definitions, and user-definable tables. You can define and maintain a single VTAM for use on multiple systems.

You can also use MVS system symbols in the TSOKEY00 parmlib member.

In addition, you can use MVS system symbols in VTAM network operator commands and Transaction Processing Facility (TPF) logon manager operator commands. Substitution text for the symbols is defined in MVS and is substituted by MVS in place of the symbols during command processing. This eliminates the need for the operator to know the exact name of a resource when issuing a command on a particular system.

You can use MVS system symbols in non-sysplex and sysplex environments.

Overview

Before coding MVS system symbols in VTAMLST, TSOKEY00, or in VTAM network operator commands or TPF logon manager operator commands, you might want to issue the DISPLAY SYMBOLS command on each system on which you plan to use the symbols. The DISPLAY SYMBOLS command is an MVS command that displays defined symbols and their substitution text for the system on which the command is issued. See [z/OS MVS System Commands](#) for information about how to use the DISPLAY SYMBOLS command.

Additional installation-defined symbols can be defined if you need them.

Consider how to best use the defined symbols and any additional symbols you might want to define. System symbols in VTAMLST are most useful for defining system identifiers and names when coding:

- Start options, especially those that specify system names and identifiers; for example, CONFIG, SSCPID, and SSCPNAME.
- A configuration list. Use the symbols in the names of the major node members contained in the configuration list.
- Definitions for major nodes and minor nodes. Use the symbols in the names of definition statements and in the values you specify on operands on the definition statements.
- Definitions for routing and dynamic reconfiguration. Use the symbols in the names of definition statements and in the values you specify on operands on the definition statements.
- The following user-definable tables:
 - APPN Class of Service (COS) definitions
 - APPN-to-subarea COS mapping table
 - Associated LU table
 - Message-flooding prevention table
 - Model name table
 - Subarea-to-APPN COS mapping table

Use the symbols in the names of macroinstructions and in the values you specify on operands on the macroinstructions.

To use MVS system symbols in VTAM network operator commands and TPF logon manager operator commands, use the symbols in the values you specify on command operands, especially those that

specify resource names (for example, ID). You can also use the static system symbols shown in [Table 3 on page 32](#).

<i>Table 3. MVS static system symbols that can be used in VTAM</i>				
Symbol name	Description	Length of text	Where defined	Default substitution text
&SYSCClone	Shorthand notation for the name of the system; often used in fields that are limited to two characters.	1–2 characters	IEASYMxx parmlib member	Last two characters of substitution text are defined to &SYSNAME symbol.
&SYSNAME	The name of the system	1–8 characters	IEASYMxx or IEASYSxx parmlib member	The processor identifier. See the z/OS MVS Initialization and Tuning Reference and the z/OS MVS Initialization and Tuning Guide for information about the processor identifier.
&SYSPLEX	The name of the sysplex	1–8 characters	COUPLExx or LOADxx parmlib member	If LOADxx does not specify the sysplex name, &SYSPLEX defaults to LOCAL until the COUPLExx member is processed.
<i>&installation-defined system symbol</i>	Up to 100 system symbols defined by your installation	1–8 characters	IEASYMxx parmlib member	None.

During VTAM startup, major node activation, resource activation, table activation, and command processing, MVS substitutes the text defined for these system symbols wherever the symbols appear in VTAMLST, TSOKEY00, and in commands.

Note: Other products, such as NCP and NetView, use VTAMLST to obtain certain definitions. See the other products' publications to determine how they support MVS system symbols in VTAMLST.

Coding guidelines

Substitution text for system symbols is defined in MVS. For complete information about how to define the symbols, including default substitution text, see the [z/OS MVS Initialization and Tuning Reference](#) and the [z/OS MVS Initialization and Tuning Guide](#). This section provides guidelines for coding the symbols in VTAMLST.

The symbols must begin with an ampersand (&) and should end with a period (.). However, if MVS finds a system symbol that does not end with a period, it substitutes text for the symbol when the next character is:

- Null (the end of text is reached)
- A character that is not alphabetical, numeric, or one of the national characters, @, #, _, or \$

If an MVS system symbol is used with a substitution value that has a length greater than the symbol name, and this substitution causes the record to overflow, then the symbol substitution will fail for the record.

For example, if &SYSCONE is defined in MVS to equal 01—SYSCONE(01)—and &SYSNAME is defined to equal HOST01A—SYSNAME(HOST01A)—, you could code the symbols in VTAMLST as follows:

```
...CONFIG=&SYSCONE . ,SSCPNAME=HOST&SYSCONE .A
```

or:

```
...CONFIG=&SYSCONE . ,SSCPNAME=&SYSNAME
```

It is recommended that you always end MVS system symbols with a period. Otherwise, you risk a higher incidence of syntax errors.

Note: For ease of reading, periods are not shown at the end of symbols in this document, except when the symbols are used in examples.

A symbol name can have one to eight characters between the ampersand and the period.

After MVS substitutes the symbols with the text defined in MVS, VTAM will ensure that the resulting values are valid. Current VTAM coding conventions still apply when MVS system symbols are used in VTAMLST. So, for example, if &SYSCONE is defined in MVS to equal 01, and CONFIG=&SYSCONE is coded in VTAMLST, VTAM accepts the resolved value, CONFIG=01. If, however, CONFIG=01&SYSCONE is coded, VTAM rejects the resolved value, CONFIG=0101, because the value specified on the CONFIG start option must be two characters in length. For information about current VTAMLST coding conventions, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Example of VTAM definitions in a sysplex environment without the use of MVS system symbols

Following is an example of the VTAM definitions required to enable VTAM to run on two MVS hosts in a sysplex without the use of MVS system symbols.

Two MVS hosts, HOST01 and HOST02, exist in a sysplex. Each host starts VTAM and is running the same VTAM application programs on its system. The following members are defined in the VTAMLST data set to enable this configuration:

```
ATCSTR01
    CONFIG=01,
    SSCPID=01,
    SSCPNAME=HOST01,
    :
ATCSTR02
    CONFIG=02,
    SSCPID=02,
    SSCPNAME=HOST02,
    :
ATCCON01
    APPLS01
    :
ATCCON02
    APPLS02
    :
APPLS01
    APPL01A  APPL ...
    APPL01B  APPL ...
    APPL01C  APPL ...
    :
APPLS02
    APPL02A  APPL ...
    APPL02B  APPL ...
    APPL02C  APPL ...
    :
```

VTAM is started from MVS HOST01 with the following command:

```
START VTAM,,, (LIST=01)
```

VTAM is started from MVS HOST02 with the following command:

```
START VTAM,,, (LIST=02)
```

Two separate start option list members are defined in VTAMLST to enable VTAM to be started on both systems:

- ATCSTR01, which specifies names and identifiers to be used when VTAM is started on HOST01
- ATCSTR02, which specifies names and identifiers to be used when VTAM is started on HOST02

Each start option list specifies the name of a configuration list to be used when VTAM is started on that system. Two configuration list members are defined in VTAMLST:

- ATCCON01, which specifies the major nodes to be activated when VTAM is started on HOST01
- ATCCON02, which specifies the major nodes to be activated when VTAM is started on HOST02

Even though identical application programs are to be activated on both systems, because they must have unique names within the sysplex, they must be defined in two separate major nodes with unique names. One major node is specified in one configuration list and the other major node is specified in the other configuration list.

Three application programs are defined in major node APPLS01 with the following names:

- APPL01A
- APPL01B
- APPL01C

Identical application programs are defined in major node APPLS02 with the following names:

- APPL02A
- APPL02B
- APPL02C

APPLS01 is specified in configuration list ATCCON01, indicating that VTAM should activate the major node defined by APPLS01 when VTAM is started on HOST01. Within member APPLS01, the APPL definition statements, APPL01A, APPL01B, and APPL01C are defined, indicating that the application programs defined by these statements should be activated when the major node defined by APPLS01 is activated.

APPLS02 is specified in configuration list ATCCON02, indicating that VTAM should activate the major node defined by APPLS02 when VTAM is started on HOST02. Within member APPLS02, the APPL definition statements APPL02A, APPL02B, and APPL02C are defined, indicating that the application programs defined by these statements should be activated when the major node defined by APPLS02 is activated.

You can see that the more systems you have in a sysplex environment, each running identical VTAM application programs, the more VTAM definitions you must code in VTAMLST.

Using MVS system symbols in a sysplex environment

Following is an example of how the definitions in the previous section can be reduced by using MVS system symbols.

The following members are defined in a VTAMLST data set that is accessible to two MVS systems, HOST01 and HOST02, in a sysplex:

```
ATCSTR00  
  
CONFIG=00,  
SSCPID=&SYSCONE.,  
SSCPNAME=&SYSNAME.,
```



```

:
ATCCON00
  APPLS
:
APPLS
  APPL&SYSCONE.A  APPL ...
  APPL&SYSCONE.B  APPL ...
  APPL&SYSCONE.C  APPL ...
:

```

How the symbols resolve when HOST01 starts VTAM:

The following substitution text is defined in MVS HOST01:

```

SYSCONE(01)
SYSNAME(HOST01)

```

VTAM is started from MVS HOST01 with the following command:

```

START VTAM,,, (LIST=00)

```

LIST=00 is not required on the START command because it is the default. VTAM substitutes every instance of &SYSCONE in VTAMLST with 01 and every instance of &SYSNAME with HOST01. As a result, VTAM uses the following definitions:

```

ATCSTR00
  CONFIG=00
  SSCPID=01,
  SSCPNAME=HOST01,
:
ATCCON00
  APPLS
:
APPLS
  APPL01A  APPL ...
  APPL01B  APPL ...
  APPL01C  APPL ...
:

```

How the symbols resolve when HOST02 starts VTAM:

The following substitution text is defined in MVS HOST02:

```

SYSCONE(02)
SYSNAME(HOST02)

```

VTAM is started from MVS HOST02 with the following command:

```

START VTAM,,, (LIST=00)

```

LIST=00 is not required on the START command because it is the default. VTAM substitutes every instance of &SYSCONE in VTAMLST with 02 and every instance of &SYSNAME with HOST02. As a result, VTAM uses the following definitions:

```

ATCSTR00
  CONFIG=00
  SSCPID=02,
  SSCPNAME=HOST02,
:
ATCCON00
  APPLS

```

```
⋮  
APPLS  
  APPL01A  APPL ...  
  APPL01B  APPL ...  
  APPL01C  APPL ...  
⋮
```

Using MVS system symbols in a non-sysplex environment

If multiple VTAMs are not in a sysplex environment, but are nonetheless related (for example, in a state university environment where one VTAM is started on campus A and one VTAM is started on campus B) one VTAMLST data set might be copied from the campus A system and installed and used to start VTAM on the campus B system. MVS system symbols can be used in this situation to reduce the system definitions in VTAMLST, much like they are used in the sysplex example "Using MVS system symbols in a sysplex environment". Instead of being two MVS hosts in a sysplex, HOST01 and HOST02 might be the two hosts in the state university environment, HOST01 on campus A and HOST02 on campus B.

Verifying a VTAM network

When you implement a VTAM network, verify that your VTAM network functions properly.

Before you begin

Make sure that the status of all resources in your network is inactive.

Procedure

To verify that your VTAM network functions properly, do the following steps:

1. Start VTAM and the appropriate trace facility (such as GTF).
2. Activate a channel-attachment major node, begin tracing, and activate lines and peripheral physical units.
3. Activate an NCP, begin tracing and intensive mode error recording, and verify that you can dump the NCP.
4. Start the NetView program, if it is installed.
5. Activate a switched major node, and begin tracing for the physical units.
6. Start the VTAM or host subsystem application programs.
7. Start VTAM traces, and activate the logical units at the peripheral physical unit.
8. Use DISPLAY commands to verify the following information:
 - Peripheral terminals and links
 - Route structure
 - Sessions between application programs and logical units
9. Activate and verify additional physical units and NCPs.
10. Simulate normal operation with all physical units.
11. Try backup and recovery procedures.
12. Halt VTAM.

What to do next

For specific information about how to perform these steps and for the syntax of operator commands, see [z/OS Communications Server: SNA Operation](#). If these steps do not yield the expected results, check the definition statements for the resources that are affected. If this does not solve the problem, see [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#).

Verifying a multiple-domain subarea network

Verify that your multi-domain subarea network functions properly after you verify a VTAM network.

Before you begin

Perform the steps [“Verifying a VTAM network” on page 36](#) for each domain in your network. After you have done this, make a list of the most important normal and backup operations in your network, and verify each of them one at a time. Then verify the network over a length of time, as in normal operation.

Procedure

To verify that your multiple-domain, subarea network functions properly, do the following steps:

1. Activate cross-domain paths and links, cross-domain resource managers (CDRMs), and cross-domain resources (CDRSCs).
2. Verify cross-domain sessions by exchanging requests across domains, between terminals and hosts, and between application programs.
3. Verify the NetView program, if it is installed.
4. Verify host backup using a channel-attached, shared NCP.
5. Verify host backup over a cross-domain link.
6. Verify automatic SSCP-SSCP session restart.
7. Halt VTAM in each domain.

Verifying a multiple-network environment

Verify that your multiple-network environment functions properly.

Before you begin

Perform the steps in both [“Verifying a VTAM network” on page 36](#) and [“Verifying a multiple-domain subarea network” on page 37](#) (in that order) for each of your networks.

Procedure

To verify that your multiple-network environment functions properly, do the following steps:

1. Activate each gateway VTAM and gateway NCP.
2. Activate major nodes required for cross-network sessions.
3. Establish required cross-network SSCP-SSCP sessions.
4. Verify cross-network sessions by exchanging requests across networks, between terminals and hosts, and between application programs.
5. Verify termination processing when an SSCP-SSCP or SSCP-gateway NCP session is deactivated.
6. Try backup and recovery procedures.
7. Verify the NetView program, if it is installed.
8. Halt VTAM in each domain in each network.

Verifying a VTAM APPN network

Verify that your VTAM APPN network functions properly.

Procedure

To verify that your VTAM APPN network functions properly, follow these steps:

1. Start VTAMs with appropriate APPN start options.
2. Activate NCP major node, XCA major node for 3172 connection, or an OSA connection.

3. Activate CP-CP session-capable APPN connections (for example, lines and PUs) to adjacent nodes and verify that CP-CP sessions are automatically started.

The following message is issued:

```
IST1096I CP-CP SESSIONS WITH adjcpname ACTIVATED
```

4. Display adjacent control points. For example, if VTAM was started with DYNADJCP=YES, issue the following command:

```
D NET,ID=ISTADJCP,E.
```

5. From VTAM network nodes, display the topology of the APPN network.
For example, issue the following command:

```
D NET,TOPO,ID=cp_name,LIST=ALL.
```

6. From VTAM end nodes, display the network node server list.
For example, issue the following command:

```
D NET,NETSRVR,E.
```

7. Terminate one of the CP-CP sessions.
For example, issue the following command:

```
V NET,TERM,LU1=cp_name,LU2=adjcp_name.
```

8. Reestablish the CP-CP session.
For example, issue the following command:

```
V NET,ACT,ID=adjcp_name,IDTYPE=CP.
```

9. To verify end node registration, first activate a resource defined with REGISTER=NETSRVR on the end node (note that application program minor nodes default to REGISTER=NETSRVR).
Then display the resource from the network node server directory database. For example, issue the following command:

```
D NET,DIRECTRY,ID=resource_name.
```

10. Verify network connectivity with any LU 6.2 resource (including control points).
For example, issue the following command:

```
D NET,APING,ID=resource_name.
```

11. From VTAM network nodes, verify that an SSCP, CP, or LU in the network can be reached, by initiating a search.
Also verify that a particular name is not duplicated in the network, which might result in session failures. For example, issue the following command:

```
D NET,DIRECTRY,ID=resource_name,SCOPE=NSARCH.
```

12. Verify sessions by exchanging requests between APPN nodes.
13. Halt VTAM. Note that the HALT or HALT,QUICK commands checkpoint the topology and routing services database.

Chapter 4. Connecting an APPN node to VTAM

This topic includes information about:

- Connections through boundary function-based transmission groups
- Multiple Connections with Parallel Transmission Groups
- Channel connections between APPN nodes
- Leased connections between APPN nodes
- IBM 3172 Nways interconnect controller connections between APPN nodes
- APPN multiple network connectivity
- Virtual-Route-Based transmission groups
- Selecting the network node server for end nodes

For a node to communicate with other nodes as a peer, rather than as a participant in a hierarchical relationship, the node must be aware of other APPN nodes. VTAM nodes use a combination of the following available definitions and information to learn about adjacent nodes:

- Definition of the physical line or port
- Definition of the adjacent link station, either defined dynamically or predefined with PU definition statements
- Transmission group number
- Adjacent control point name

Before and during link activation between adjacent link stations, nodes also obtain information about the capability of an adjacent node through a message unit called an exchange identification (XID), which is used to convey node and link characteristics. XIDs are exchanged by adjacent nodes to establish and negotiate link and node characteristics. After link activation, XIDs are used to communicate changes in these characteristics and to negotiate the TG number used to represent the connection.

APPN nodes are connected using type 2.1 connections. APPN connections provide the same benefits as LEN connections. With APPN connections, you do not need to define CDRMs, PATH definitions, or ADJSSCP tables. You can dynamically define switched connections and independent logical units, and the connectivity of applications and dependent LUs is not limited.

Notes:

1. If you use VR-based transmission groups, CDRMs, PATH definitions, and ADJSSCP tables must be considered.
2. When changing the node roles of nodes in the network, a node should be taken down using normal takedown procedures. Forcing a node down and then bringing it up with a different node role can cause errors in the topology database, because error recovery might still be in progress when the node is brought back up.

Connections through boundary function-based transmission groups

Nodes are connected by transmission groups (TGs). A TG connects a pair of adjacent link stations across the transmission medium. A boundary function (BF)-based APPN TG is used to connect two APPN nodes that use the NCP or VTAM boundary function.

Transmission groups are described with the following standard APPN characteristics:

- Cost per connect time
- Cost per byte

- Security
- Propagation delay
- Effective capacity
- User-defined values

The values specified for these characteristics are saved in the topology database and are used to determine the weight of the TG for route calculation. The values you specify represent the characteristics of the link; the values do not affect the actual physical characteristics of the link. If the physical characteristics of the link change, the values specified for the TG characteristics should probably be changed to reflect the new physical characteristics. For more information about the topology database, see [“Topology database” on page 14](#).

By default, all TGs assume the same characteristics. To distinguish one TG from another, TG characteristics can be specified in the following ways:

- As TG profiles, which are groups of TG characteristics that can be applied to several PUs
- On definition statements

The characteristics of TGs owned by other nodes are learned through topology database updates (TDUs).

Multiple connections with parallel transmission groups

Parallel TGs for multiple connections between APPN nodes are supported. This support means that a node can have multiple links to the same adjacent node. Parallel TGs have different TG numbers and can have different TG characteristics assigned to them. Multiple-link TGs (many links with one TG number) are not supported in APPN by VTAM. The support for parallel TGs matches that of most other APPN products.

If you have parallel TGs with CONNTYPE=APPN coded on the adjacent link stations, you can establish one or more APPN connections or one LEN connection between the two nodes. You cannot combine LEN and APPN connections or establish more than one LEN connection to the same adjacent node over APPN-capable TGs.

Support for multiple connections allows an independent LU in the APPN network to establish sessions through or into the subarea network by using several boundary-function TGs. An LU can have multiple sessions with one or more LUs through multiple boundary-function TGs and through different physical paths that function as parallel TGs. This support is also available for TGs attached to LEN nodes.

When coding CONNTYPE=APPN (or letting it default to that value), XID exchange rules for parallel TGs are enforced. This means that if you have PUs with duplicate CPNAMES, you should resolve this duplication before that node becomes an APPN node. If the node is to remain a LEN node, then CONNTYPE=LEN should be coded on the PU statement. Coding CONNTYPE=LEN results in VTAM avoiding the test for duplicate CPNAMES.

Channel connections between APPN nodes

Type 2.1 channel connections can be used to connect adjacent APPN nodes. Adjacent APPN nodes can be other IBM Z hosts or other devices that support APPN channel connections.

This section contains information about the following types of channel connections:

- Multipath channel connections
- Composite network node channel connections

Multipath channel connections

Multipath channel (MPC) connections allow you to code a single transmission group (TG) that uses multiple write-direction and read-direction subchannels. The subchannels may be assigned to one or more physical channels based on how you have defined your I/O subsystem to the operating system. Because each subchannel operates in only one direction, the half-duplex turnaround time that occurs with other channel-to-channel connections is reduced.

Levels of MPC capability provided by VTAM

VTAM provides the following levels of MPC capability:

- High-performance data transfer (HPDT)

HPDT MPC connections provide more efficient transfer of data than non-HPDT MPC connections. They do this by using HPDT services to provide the following functions:

- Data packing without data movement: This process decreases consumption of CPU cycles by reducing the internal movement of data, thus increasing the availability of MIPS (million instructions per second) for user processing.
- Chain scheduling of channel programs: This process reduces operating system I/O invocations and CPU overhead.

HPDT MPC can also run in HPDT Packing mode. HPDT Packing reduces data stream fragmentation. See [“HPDT Packing” on page 45](#) for more information.

HPDT services are available over connections to other nodes that implement HPDT MPC. Applications that use high-performance data transfer services rely on HPDT MPC connections for path length reductions and performance enhancements when sending and receiving data. For more information about high-performance data transfer services, see the [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

HPDT MPC uses the services of the communications storage manager (CSM). CSM is a storage facility that allows authorized host applications to set up data in buffers that can be addressed and accessed by other authorized host applications, eliminating the need to copy data as it is being prepared to be sent. CSM is installed with VTAM. Maximum CSM storage limits are defined in the CSM parmlib member, IVTPRM00. For more information about CSM, see [“Communications storage manager \(CSM\)” on page 340](#).

- Non-HPDT

Non-HPDT MPC connections do not use HPDT services. Non-HPDT MPC connections can be considered synonymous with APPN host-to-host channel (AHHC) connections. See [“Using MPC connections in an APPN network” on page 41](#).

- Subarea

Subarea MPC connections are described in [“Multipath channel connections” on page 89](#).

Using MPC connections in an APPN network

You can use MPC connections in an APPN network to connect VTAM to adjacent APPN nodes.

How VTAM uses HPDT MPC connections in an APPN network

HPDT MPC is used by VTAM in an APPN network to enable the following connections:

- APPN node-to-node connections (ANNCs)

ANNC connections can be between a VTAM APPN host node and an adjacent VTAM APPN host node or between a VTAM APPN host node and an adjacent nonhost APPN node.

Connections between two VTAM APPN host nodes in a sysplex can also use the cross-system coupling facility (XCF). See [“Dynamic definition of VTAM-to-VTAM connections” on page 350](#) for more information.

How VTAM uses non-HPDT MPC connections in an APPN network

Non-HPDT MPC is used by VTAM to enable connections only between a VTAM APPN host node and an adjacent VTAM APPN host node. These connections are called APPN host-to-host channel (AHHC) connections. AHHC connections can be considered synonymous with non-HPDT MPC connections.

Steps involved in using HPDT and non-HPDT MPC connections in an APPN network

To use an MPC connection in an APPN network, you must complete the following steps:

1. Define the MPC connection.
2. Define the partner node.
3. Activate the resources that enable the MPC connection.

Note: If VTAM is using the functions of XCF in a sysplex, these steps can be performed dynamically. See [“Dynamic definition of VTAM-to-VTAM connections” on page 350](#) for more information.

Defining an MPC connection:

You define an MPC connection in the transport resource list (TRL) major node. The characteristics of the connection are defined on a TRLE definition statement. You must code a TRLE definition statement for each MPC connection.

Defining the partner node:

As explained earlier, VTAM uses an MPC connection to connect to:

- An adjacent APPN node

You define an adjacent APPN node (host or nonhost) in the local SNA major node. The TRLE operand on the PU definition statement in the local SNA major node specifies the name of the TRLE definition statement in the TRL major node. The TRL major node defines the MPC connection to be used between VTAM and the adjacent APPN node.

For TCP/IP to use an MPC connection, the device name on the DEVICE statement must match the MPC TRLE name (for MPCPTP devices). For more information, see the [z/OS Communications Server: IP Configuration Guide](#) and the [z/OS Communications Server: IP Configuration Reference](#).

Activating the resources that enable the MPC connection:

After creating the TRL major node and the major node that defines the partner node (local SNA or XCA), you can use the MPC connection by activating the appropriate resources. For an MPC connection between VTAM and an adjacent APPN node, follow the steps below:

1. Activate the TRL major node.
2. Activate the local SNA major node.
3. Activate the PU that defines the connection to the adjacent node as an APPN PU.

Multiple TRL major nodes are allowed to exist in VTAM, and each can be dynamically modified using V NET,ACT,ID= with the UPDATE operand.

If UPDATE=ALL is specified, the VTAMLST member designated with the ID operand replaces the existing TRL major node. However, note that TRLEs in use are not deleted and remain in the TRL major node along with the new TRLEs.

Note: When there is only one active TRL major node, UPDATE=ALL deletes or replaces all inactive TRLEs. When there are multiple active TRL major nodes, UPDATE=ALL deletes or replaces only the inactive TRLEs that exist in the TRL major node being activated. If an inactive TRLE is defined in another TRL major node, it will not be deleted and a name conflict will result if an attempt is made to replace it later.

If UPDATE=ADD is specified, the TRLE entries in the VTAMLST member identified on the ID operand are added to the TRL major node if they do not already exist in any TRL major node.

MPC connection example

Figure 6 on page 43 shows two VTAM network nodes connected through an MPC connection. Corresponding sample definitions follow [Figure 6 on page 43](#).




Figure 6. MPC connection between two VTAM network nodes

Definitions for network node 1	Definitions for network node 2
TRL Major Node TRLNN1 VBUILD TYPE=TRL MPCLN1 TRLE LNCTL=MPC, MAXBFRU=16, READ=(BCA), WRITE=(BCB) Local SNA Major Node LSNA1 VBUILD TYPE=LOCAL LSNAPU1 PU TRLE=MPCLN1, ISTATUS=ACTIVE, CONNTYPE=APPN, CPCP=YES	TRL Major Node TRLNN2 VBUILD TYPE=TRL MPCLN2 TRLE LNCTL=MPC, MAXBFRU=16, READ=(BCB), WRITE=(BCA) Local SNA Major Node LSNA2 VBUILD TYPE=LOCAL LSNAPU2 PU TRLE=MPCLN2, ISTATUS=ACTIVE, CONNTYPE=APPN, CPCP=YES

How the level of MPC is determined

The level of MPC used for an MPC connection is automatically determined by a negotiation between the partner nodes. You can control the level used by coding the MPCLEVEL operand on the TRLE definition statement that defines the MPC connection.

When HPDT MPC is used

If both partner nodes support HPDT MPC, HPDT MPC is automatically used.

VTAM supports HPDT MPC if it is defined as an HPR APPN node that provides RTP-level HPR support. See [“High-Performance Routing \(HPR\)” on page 388](#) for information about how VTAM is defined as an HPR APPN node that provides RTP-level HPR support.

In [Figure 6 on page 43](#), assume that both VTAM network nodes are nodes higher than VTAM Version 4 Release 4, each providing RTP-level HPR support. The MPC connection between them, therefore, is an HPDT MPC connection.

When Non-HPDT MPC is used

If either partner node does not support HPDT MPC, non-HPDT MPC is automatically used.

Keep in mind that VTAMs before Version 4 Release 4 do not support HPDT MPC, and therefore, non-HPDT MPC will be used.

In [Figure 6 on page 43](#), assume that VTAM network node 1 is a VTAM Version 4 Release 4 or higher node that does not provide RTP-level HPR support and that VTAM network node 2 is a VTAM node before Version 4 Release 4. Because VTAM releases before Version 4 Release 4 do not support HPDT MPC, the MPC connection between them is a non-HPDT MPC connection.

Of course, if both nodes are VTAM Version 4 Release 4 or lower network nodes, the MPC connection between them would also be a non-HPDT MPC connection.

Controlling the level of MPC used

Use the MPCLEVEL operand on the TRLE definition statement to control the level of MPC used for the connection. When VTAM and the adjacent node support HPDT MPC, the connection automatically uses HPDT MPC, unless VTAM includes MPCLEVEL=NOHPDT on its TRLE definition statement. When the adjacent node is another VTAM node that supports HPDT MPC, if MPCLEVEL=NOHPDT is coded on the TRLE definition statement of at least one of the VTAM nodes, the connection uses non-HPDT MPC.

Note: There are situations in which you must code MPCLEVEL=NOHPDT. Otherwise, the connection or session requests over the connection will fail. These situations include:

- Both nodes are VTAM Version 4 Release 4 or higher nodes and one of the VTAM nodes does not provide RTP-level HPR support. In this case, MPCLEVEL=NOHPDT must be specified on the TRLE definition statement in the VTAM node that does not provide RTP-level HPR support or the connection attempt will fail. See the following example for a full explanation of this situation.
- Both nodes are VTAM Version 4 Release 4 nodes that provide RTP-level HPR support, but the HPR operand on the PU definition statement in the local SNA major node in at least one of the VTAM nodes is coded as follows:
 - HPR=NO
 - HPR=YES, when the HPR start option specifies HPR=(RTP,NONE) or HPR=(RTP, ANR)

In these cases, MPCLEVEL=NOHPDT must be specified on the TRLE definition statement in at least one of the VTAM nodes or the connection attempt will fail.

In Figure 6 on page 43, assume that both VTAM network nodes are VTAM nodes, but that VTAM network node 1 provides RTP-level HPR support and VTAM network node 2 does not. Because the default value for the MPCLEVEL operand is HPDT, during negotiation, both VTAM nodes agree to use an HPDT MPC connection. But, because VTAM network node 2 does not provide RTP-level HPR support, the connection attempt fails. For the connection to be successfully established, it must be a non-HPDT MPC connection. You must code MPCLEVEL=NOHPDT on the TRLE definition statement that defines the connection in VTAM network node 2 to make this a non-HPDT MPC connection, as follows:

Definitions for network node 1 HPR=RTP	Definitions for network node 2 HPR=NONE
<pre>TRL Major Node TRLNN1 VBUILD TYPE=TRL MPCLN1 TRLE LNCTL=MPC, MAXBFRU=16, READ=(BCA), WRITE=(BCB) Local SNA Major Node LSNA1 VBUILD TYPE=LOCAL LSNAPU1 PU TRLE=MPCLN1, ISTATUS=ACTIVE, CONNTYPE=APPN, CPCP=YES</pre>	<pre>TRL Major Node TRLNN2 VBUILD TYPE=TRL MPCLN2 TRLE LNCTL=MPC, MAXBFRU=16, READ=(BCB), WRITE=(BCA), MPCLEVEL=NOHPDT Local SNA Major Node LSNA2 VBUILD TYPE=LOCAL LSNAPU2 PU TRLE=MPCLN2, ISTATUS=ACTIVE, CONNTYPE=APPN, CPCP=YES</pre>

MPC dynamics

With MPC dynamics support, if at least one WRITE and one READ subchannel path is allocated successfully, an MPC connection is activated. Additional paths (defined but not ONLINE) in an MPC group can later be dynamically added to the active group using the MVS command `VARY device ONLINE`. For example, if there is a need for an increase in capacity to allow for extra traffic over a channel, additional paths can be added to the active group without disruption. Similarly, paths can be deleted from the active group when no longer needed using the MVS command `VARY device OFFLINE`.

Note: By default, the last READ and last WRITE subchannel paths in an MPC group cannot be deleted. To be able to use `VARY device OFFLINE` to delete the last READ or last WRITE subchannel path in an MPC group, include `LASTRW=ALLOW` on the TRLE definition statement for the MPC connection. For details on the TRLE definition statement, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

If a path fails, it can be dynamically added back to the active group without disruption [automatically using the Event Notification Facility (ENF) channel recovery support or by issuing the MVS command `VARY device ONLINE`].

When an MPC connection is activated, if the minimum number of WRITE and READ subchannel paths cannot be allocated successfully, either because they are offline or no valid path exists to the connecting host, then activation is suspended until the minimum number becomes available. This is true unless the start option specification `MPCACT=NOWAIT` is in effect.

HPDT Packing

Restriction: HPDT Packing can be enabled only for non-XCF point-to-point connections. Therefore XCF and OSA devices are excluded.

The HPDT device driver was initially designed for large data transfer and within the architectural limitations of the I/O subsystem. HPDT is efficient at transporting large pieces of data but relatively inefficient when building a data stream with multiple small to medium size packets (packets from 257 to 2KB bytes, including headers). A write data stream built from this type of workload results in a significant number of alignment bytes. These alignment bytes are transmitted across the media but are not used by the receiver. In effect, the alignment bytes significantly reduce the media bandwidth.

HPDT Packing can be enabled to eliminate most of these alignment bytes and effectively increase the media bandwidth. However, HPDT Packing is a compromise between CPU cycles, storage, and channel bandwidth. There are two side effects of enabling HPDT Packing:

- CPU use increases because of the cost of moving the data into the transmit packing buffer.
- Storage use increases because fixed packing transmit buffers are allocated.

Recommendation: In cases where the HPDT MPC connection is to a router or through a channel extender (and packet size is small to medium), you should consider enabling HPDT Packing. In these cases, the benefits of a densely packed data stream might exceed the costs in CPU and storage resources. If channel bandwidth is not a concern, enabling HPDT Packing is not recommended.

Determining the storage cost of enabling HPDT Packing

HPDT Packing is enabled on a TRLE basis. When enabled, seven packing buffers are acquired for each write device defined in the TRLE. Packing buffer size does not vary within the TRLE. Packing buffers are obtained from the smallest CSM data space pool that can contain the entire buffer.

Determining packing buffer size

The packing buffer size is equal to the number of 4 KB pages (minus 1 page) specified by the `MAXBFRU` parameter on the TRLE definition in the adjacent link station (or equivalent when the adjacent link station is a router). This value can be determined without accessing the adjacent link station when the TRLE is active. Issue the `DISPLAY NET,ID=trlename` command for the TRLE for which you are considering enabling HPDT Packing. The `MAXBFRU` value of the adjacent link station is contained in the response. The response looks similar to the following example:

```
IST097I DISPLAY ACCEPTED
IST075I NAME = TRLE1A, TYPE = TRLE
IST486I STATUS= ACTIV---E, DESIRED STATE= ACTIV
IST087I TYPE = LEASED, CONTROL = MPC, HPDT = YES
IST1715I MPCLEVEL = HPDT MPCUSAGE = SHARE
IST1717I ULPID = AHHCPU1
IST1801I UNITS OF WORK FOR NCB AT ADDRESS X'02EAE018'
IST1802I CURRENT = 0 AVERAGE = 1 MAXIMUM = 2
>IST1577I HEADER SIZE = 4092 DATA SIZE = 16 STORAGE = ***NA***
IST1221I WRITE DEV = 0508 STATUS = ACTIVE STATE = ONLINE
IST1577I HEADER SIZE = 4092 DATA SIZE = 16 STORAGE = DATASPACE
IST1221I READ DEV = 0408 STATUS = ACTIVE STATE = ONLINE
IST314I END
```

The DATA SIZE value reported in the first IST1577I message reflects the MAXBFRU parameter in the adjacent link station. In this case, the adjacent link station defined a data stream maximum size of 16 pages. Therefore, if HPDT Packing is enabled for this TRLE, the packing buffer size will be 15 pages, or 60 KB. Because 7 packing buffers are allocated for the single write device, the total storage cost in CSM fixed data space is 7*60 KB or 420 KB.

Packing buffer size versus CSM data space pool size

The packing buffers are obtained from the smallest CSM data space pool that can contain the entire buffer. If the packing buffer size does not equal one of the CSM data space pool sizes, storage is wasted.

In the example above, the packing buffer size is exactly equal to the second largest CSM data space pool size; therefore no storage is wasted. The worst case occurs when DATA SIZE is equal to 10, meaning 9 page (36 KB) packing buffers are used. Because the packing buffers cannot be acquired from the 32 KB pool, they are acquired from the 60 KB pool. The last 24 KB of each packing buffer is wasted. 7*24 KB equals 160 KB of wasted storage for a single write device.

The previous example explains why you should adjust the MAXBFRU value in the adjacent node (if necessary) before enabling HPDT Packing, so that the packing buffer size is equal to one of the CSM data space pools sizes.

Table 4 on page 46 shows the correlation between DATA SIZE (MAXBFRU in the adjacent link station), the packing buffer size, the CSM data space pool size from which the packing buffers are allocated, and any apparent storage waste because of mismatch of packing buffer size and CSM data space pool size.

<i>Table 4. HPDT Packing - packing buffer size, CSM pools size, and waste per packing buffer</i>			
DATA SIZE	Packing buffer size	CSM pool used	Waste per packing buffer
2	4 KB	4 KB	None
3	8 KB	16 KB	8 KB
4	12 KB	16 KB	4 KB
5	16 KB	16 KB	None
6	20 KB	32 KB	12 KB
7	24 KB	32 KB	8 KB
8	28 KB	32 KB	4 KB
9	32 KB	32 KB	None
10	36 KB	60 KB	24 KB
11	40 KB	60 KB	20 KB
12	44 KB	60 KB	16 KB
13	48 KB	60 KB	12 KB
14	52 KB	60 KB	8 KB
15	56 KB	60 KB	4 KB
16	60 KB	60 KB	None

Composite network node channel connections

Figure 7 on page 47 shows how to connect an NCP in a composite network node to a network node.

Note: This example does not include all definitions needed for an NCP major node.




Figure 7. Type 2.1 channel connection between a composite network node and a network node

Composite network node definitions

Define the NCP major node as follows:

```

:
GRP3AAA1 GROUP LNCTL=CA, CA=TYPE5, NCPCA=ACTIVE
LN3AAA1 LINE ADDRESS=04, TIMEOUT=840.0, CASDL=420.0, TRANSFR=254, X
INBFRS=128, ANS=CONT
P3A21AA1 PU PUTYPE=2, XID=YES, X
CONNTYPE=APPN, CPCP=YES
:

```

Network node definitions

Define the local SNA major node as follows:

```

LSNA3AA VBUILD TYPE=LOCAL
LSNA3APA PU PUTYPE=2, CUADDR=050, XID=YES, *
MAXBFRU=15, CONNTYPE=APPN, *
CPCP=YES

```

Leased connections between APPN nodes

A leased connection can be used to connect a composite network node to a composite network node or to other nodes that support leased connections. [Figure 8 on page 48](#) shows a leased connection.




Figure 8. Leased connection between two composite network nodes

The following examples show how to connect two composite network nodes.

Connect the NCP major node for the first host as follows:

```
GRP4A6  GROUP LNCTL=SDLC,DIAL=NO,MODE=SEC,TYPE=NCP,          X
          CLOCKNG=EXT,DUPLEX=HALF,SPEED=1200
*
LN4A6   LINE ADDRESS=04E,TADDR=C1
*
          SERVICE ORDER=(P4A4956C)
*
P4A4956C PU PUTYPE=2,XID=YES,                                X
          MAXDATA=256,MAXOUT=1,PASSLIM=1,                     X
          ANS=CONTINUE,CONNTYPE=APPN,CPCP=YES
```

Connect the NCP major node for the second host as follows:


```
GRP3A9  GROUP LNCTL=SDLC,DIAL=NO,TYPE=NCP,                  X
          CLOCKNG=EXT,DUPLEX=HALF,SPEED=1200
*
LN3A11  LINE ADDRESS=03F
*
          SERVICE ORDER=(P3A4956M)
*
P3A4956M PU PUTYPE=2,ADDR=C1,XID=YES,                        X
          MAXDATA=256,MAXOUT=1,PASSLIM=1,                     X
          ANS=CONTINUE,CONNTYPE=APPN,CPCP=YES
```

IBM 3172 Nways Interconnect Controller connections between APPN nodes

You can use an IBM 3172 Nways Interconnect Controller connection to connect a composite network node, end node, or network node with any other type of APPN node.

Note: VTAM uses single-route broadcasts (rather than all-routes broadcasts) when attempting to connect to nodes on token-ring networks attached through the IBM 3172 Nways Interconnect Controller. When a node can be reached only through bridges, there must be a route from the IBM 3172 Nways Interconnect Controller to the node, and that route must traverse bridges that are configured to route single-route broadcasts. See the *Token-Ring Network Architecture Reference* for more information about all-routes and single-route broadcasts.

Figure 9 on page 49 shows how to connect two network nodes using IBM Nways Interconnect Controller.



APPN Connection -.-

Figure 9. Two network nodes connected using an IBM 3172 Nways Interconnect Controller

Examples

As shown in the following example, by coding DYNPU in the XCA major node, network node 1 does not need a switched major node; however, network node 2 must initiate the session using the DIALNO coded in its switched major node.

XCA major node in network node 1

```
XCA1A1L  VBUILD  TYPE=XCA
PORT1A1L  PORT    MEDIUM=RING, ADAPNO=1, SAPADDR=4, CUADDR=500, TIMER=254
GP1A1L1   GROUP   DIAL=YES, ANSWER=ON, CALL=INOUT,                X
                                     DYNPU=YES, DYNPUPFX=BA
L1A1L1AA  LINE
P1A1L1AA  PU
```

XCA major node in network node 2

```
XCAA1L  VBUILD  TYPE=XCA
PORTAA1L  PORT    MEDIUM=RING, ADAPNO=1, SAPADDR=4, CUADDR=503, TIMER=254
GPAA1L1   GROUP   DIAL=YES, ANSWER=ON, CALL=INOUT
LAA1L11A  LINE
PAA1L11A  PU
```

Switched major node in network node 2

```
SWXCAA  VBUILD  TYPE=SWNET
SWP1A1  PU      MAXDATA=256, ADDR=11, CPNAME=SSCP1A,                X
                                     PUTYPE=2
PTH1A1L  PATH    DIALNO=1A04003A11111111,                          X
                                     GRPNM=GPAA1L1
```

Note: Asynchronous transfer mode (ATM) networks accessed through LAN emulation appear to VTAM to be Ethernet or Ethernet-type LANs or token-ring networks and are defined to VTAM as such. ATM networks accessed through native ATM are defined to VTAM differently than those accessed through

LAN emulation. See “ATM native connections” on page 55 for information about defining ATM native connections.

Using a connection network

A connection network is a representation of a shared access transport facility (SATF), such as a local area network (LAN). This arrangement enables nodes identifying their connectivity to the SATF by a common virtual routing node to communicate without having individually defined connections to one another. This is illustrated in the following figures.

In Figure 10 on page 50, when an LU-LU session between resources at AS400A and HOSTB is set up, the optimal route is directly through the token ring between the two nodes. However, if AS400A and HOSTB are not directly defined to each other or are not connected to the same connection network, the indirect route is selected through HOSTA.

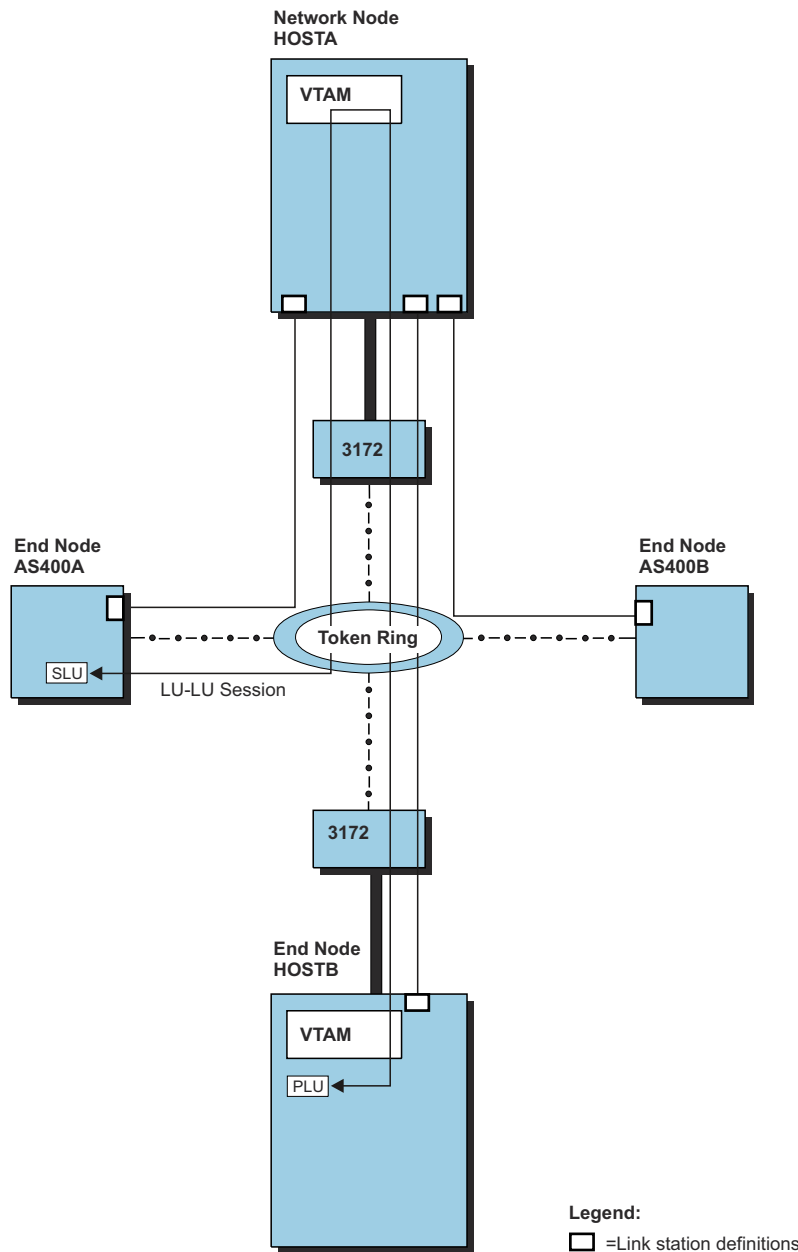


Figure 10. VTAM attachment to a LAN—No meshed connection definitions

Optimal route calculation is achieved in one of two ways:

- Meshed connection definitions

- Connection network definition

A definition for the connection between each pair of nodes, $(n*(n-1))$ definitions where n equals the number of nodes, enables optimal route calculation between all nodes on the SATF. In [Figure 11](#) on page 51, for example, the pictured LU-LU session no longer traverses HOSTA.




Figure 11. VTAM attachment to a LAN—Meshed connection definitions provide optimal route calculation

In a large network, however, this system definition is extensive. System definition is greatly reduced by defining a connection network to represent the shared access transport facility (here, the token ring). In a connection network, end nodes need to only define a connection to a virtual node, as shown in [Figure 12](#) on page 52.

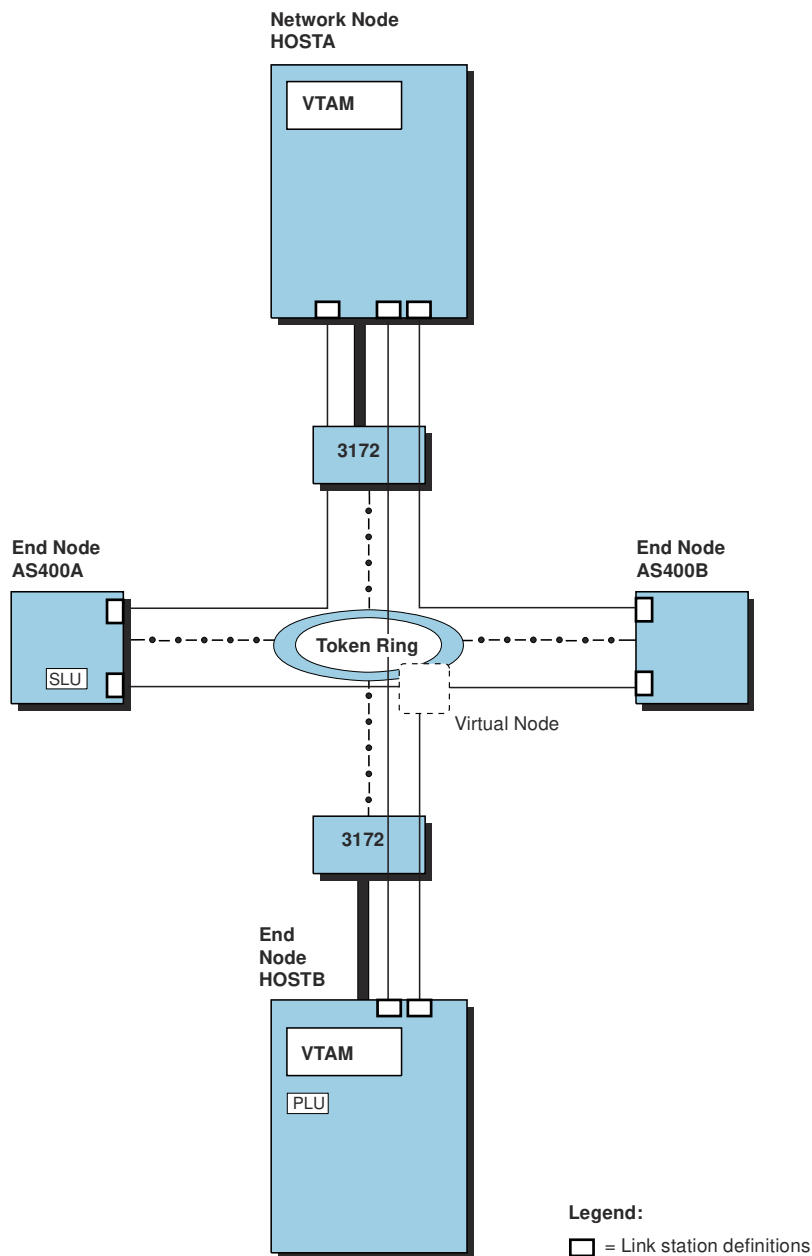


Figure 12. VTAM attachment to a connection network reduces required connection definitions (token ring)

The virtual node is reported to the topology database and can be chosen as the intermediate node during route calculation. As shown in Figure 13 on page 53, BINDs can then be routed directly to the destination node over a dynamically created TG. Thus, the definition of a connection network to represent the token ring reduces required system definition and enables optimal route calculation.

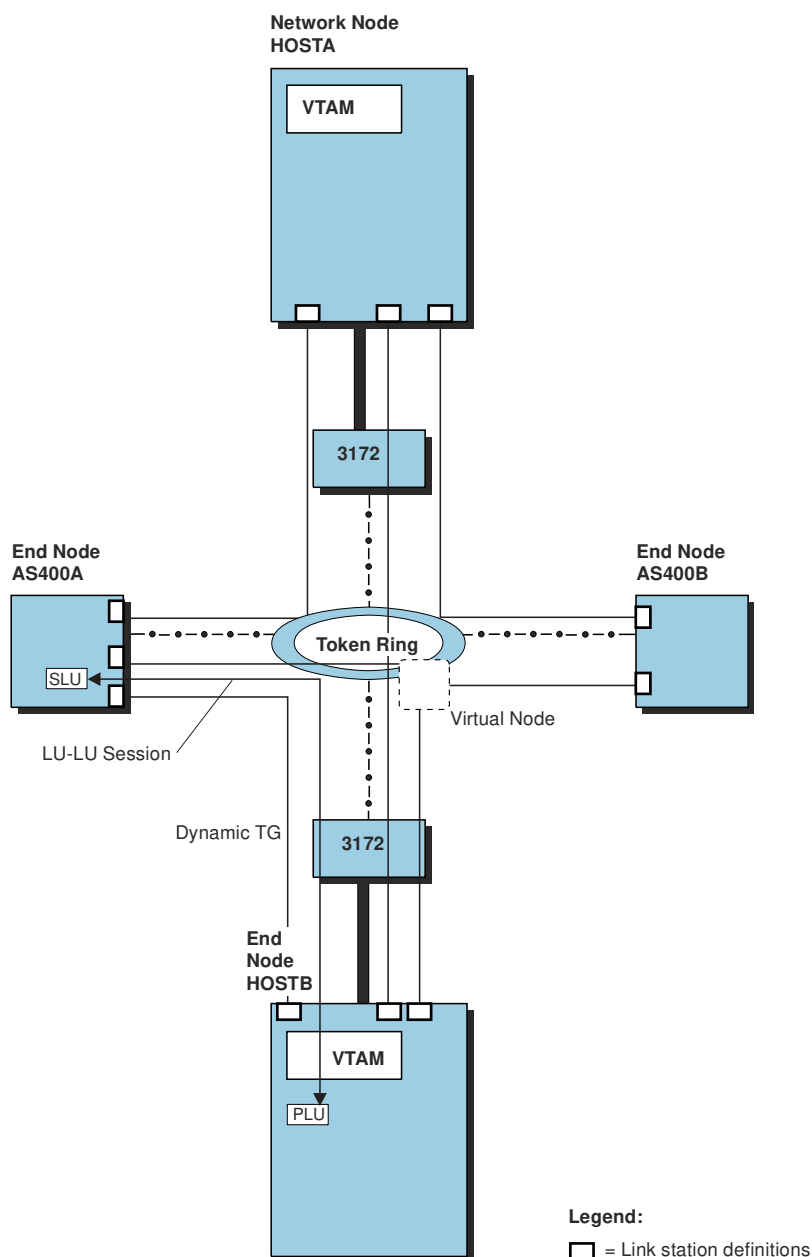


Figure 13. VTAM attachment to a connection network also enables optimal route calculation (token ring)

To define a connection to a connection network, include the following in the NCP major node, the XCA major node, or the LAN major node:

- In the NCP major node [NTRI (NCP token ring interconnect) physical lines]:
 - Code the VNNAME and VNGROUP operands on the LINE definition statement. The VNNAME operand specifies the CP name of the virtual node. The VNGROUP operand specifies the name of the logical group containing dial-out links through the connection network named on the VNNAME operand.
 - Note:** DYNPU=YES is the default when VNNAME and VNGROUP are coded. For more information about DYNPU, see Chapter 8, “Defining resources dynamically,” on page 167.
 - Optional operands on the LINE definition statement include the TGP operand to indicate a TG profile containing TG characteristics, or operands for individual TG characteristics.
- In the XCA major node (IBM 3172 Nways Interconnect Controller lines):

- Code the VNNAME and VNGROUP operands on the PORT definition statement. The VNGROUP operand specifies the GROUP in this major node containing the lines for the virtual node named on the VNNAME operand.

Note: DYNPU=YES is the default when VNNAME and VNGROUP are coded. DIAL=YES is required on the GROUP named on the VNGROUP operand.

- Optional operands on the PORT definition statement include the TGP operand to indicate a TG profile containing TG characteristics, or operands for individual TG characteristics.

- In the LAN major node:

- Code the VNNAME and VNGROUP operands on the PORT definition statement. The VNGROUP operand specifies the GROUP in this major node containing the lines for the virtual node named on the VNNAME operand.

Note: DYNPU=YES is the default when VNNAME and VNGROUP are coded. DIAL=YES is required on the GROUP named on the VNGROUP operand.

- Optional operands on the PORT definition statement include the TGP operand to indicate a TG profile containing TG characteristics, or operands for individual TG characteristics.

Notes:

1. NCP Version 7 Release 1 is required for NCP/Token-Ring interconnection (NTRI) support of connection network.
2. If CP-CP sessions are required between two nodes on the shared access transport facility, the dialing-out node must define the PU for any node it is to call; routes traversing a virtual node cannot be used for CP-CP sessions.
3. The dial-out node can use only a dynamically defined PU for connectivity. The dial-in node attempts to use a predefined PU, if one exists. Otherwise, it uses a dynamically defined PU.
4. An ADJCP definition statement with the VN=YES operand can optionally be placed into the adjacent control point major node. The VN=YES operand indicates that this adjacent CP is a virtual node. The VN operand cannot be specified when the NN operand or the DYNLU=YES operand is also specified.
5. Both VTAM end nodes and network nodes can define a connection to a virtual node.
6. A connection network can be a node in an HPR route.
7. Available lines must exist for the call through the connection network both within the dial-out node and the dial-in node. For the dial-out node, available lines must exist in the group specified by the VNGROUP keyword. The dial will fail if no connectable line is found. With multiple virtual node definitions between two real nodes, if the VNGROUP keywords specify different groups, because either route through the virtual node may be selected (depending upon the topological weight), it causes the connection network function to attempt to locate a connectable line in a group without one. For the dial-in node, a connectable line must be available for the call through the connection network. If all lines are busy with other calls, the call through the connection network will fail.
8. By default, dynamic connection network PUs take a name in the format *CNVxxxxx*, where *xxxxx* is a unique value that is generated and concatenated to the default *CNV* prefix to create the eight-character PU name. To specify a different prefix (two characters), use the DYNVNPFx start option.
9. A model PU definition can be created to customize the characteristics of dynamically created PUs. Use the DYNTYPE=VN operand for the model PU definition in the model major node.

IBM Open Systems Adapter connections between APPN nodes

VTAM can connect to the following shared access transport facilities (SATFs) through the IBM Open Systems Adapter:

- Local area network (LAN) connections
- Asynchronous transfer mode (ATM) networks accessed through LAN emulation or native ATM

LAN connections

VTAM supports the following types of LANs through the IBM Open Systems Adapter (configured in SNA mode):

- Ethernet or Ethernet-type LAN
- Token-ring network
- Fiber distributed data interface (FDDI)

LANs attached to VTAM through an IBM Open Systems Adapter are defined to VTAM in the same way as LANs attached through an IBM 3172 Nways Interconnect Controller. To define LAN connections through the IBM Open Systems Adapter, apply the information about defining LAN connections through the IBM 3172 Nways Interconnect Controller, found in [“IBM 3172 Nways Interconnect Controller connections between APPN nodes”](#) on page 48.

ATM connections

ATM is a switching technology that provides fast, reliable, simultaneous transfer of data, voice, and video. VTAM supports ATM technology by enabling communication across ATM networks—both public (wide area networks, or WANs) and private (campus networks)—accessed through:

- LAN emulation
- Native ATM

ATM LAN emulation connections

ATM networks accessed through LAN emulation are attached to VTAM through an IBM Open Systems Adapter (configured in SNA mode). They appear to VTAM as though they are Ethernet or Ethernet-type LANs or token-ring networks. ATM LAN emulation connections offer the benefits of fast data transfer, but they do not offer the full services of ATM networks, such as guaranteed bandwidth and Quality of Service.

ATM LAN emulation connections through an IBM Open Systems Adapter are defined to VTAM in the same way as LAN connections through an IBM 3172 Nways Interconnect Controller. To define ATM LAN emulation connections through the IBM Open Systems Adapter, apply the information about defining LAN connections through the IBM 3172 Nways Interconnect Controller, found in [“IBM 3172 Nways Interconnect Controller connections between APPN nodes”](#) on page 48.

ATM native connections

ATM networks accessed through native ATM are also attached to VTAM through an IBM Open Systems Adapter (configured in HPDT ATM native mode). They offer the full services of ATM networks, such as guaranteed bandwidth and Quality of Service.

VTAM enables multiprotocol attachments to ATM networks (both public and private) accessed through native ATM. These multiprotocol attachments are achieved through a high-performance data transfer (HPDT) multipath channel (MPC) connection to an IBM Open Systems Adapter. A single HPDT MPC connection can be simultaneously shared by multiple products using different communication protocols, provided those other products also support HPDT MPC connections. For example, APPN and non-APPN data can be simultaneously sent and received across an ATM network over the same HPDT MPC connection. For more information about HPDT MPC connections, see [“Multipath channel connections”](#) on page 40.

Note: VTAM can route only HPR APPN data across ATM native connections. Routing of subarea and non-HPR APPN data across ATM networks requires ATM LAN emulation connections.

Defining ATM native connections to VTAM

Figure 14 on page 56 shows a basic ATM configuration enabling HPR APPN communication through native access.

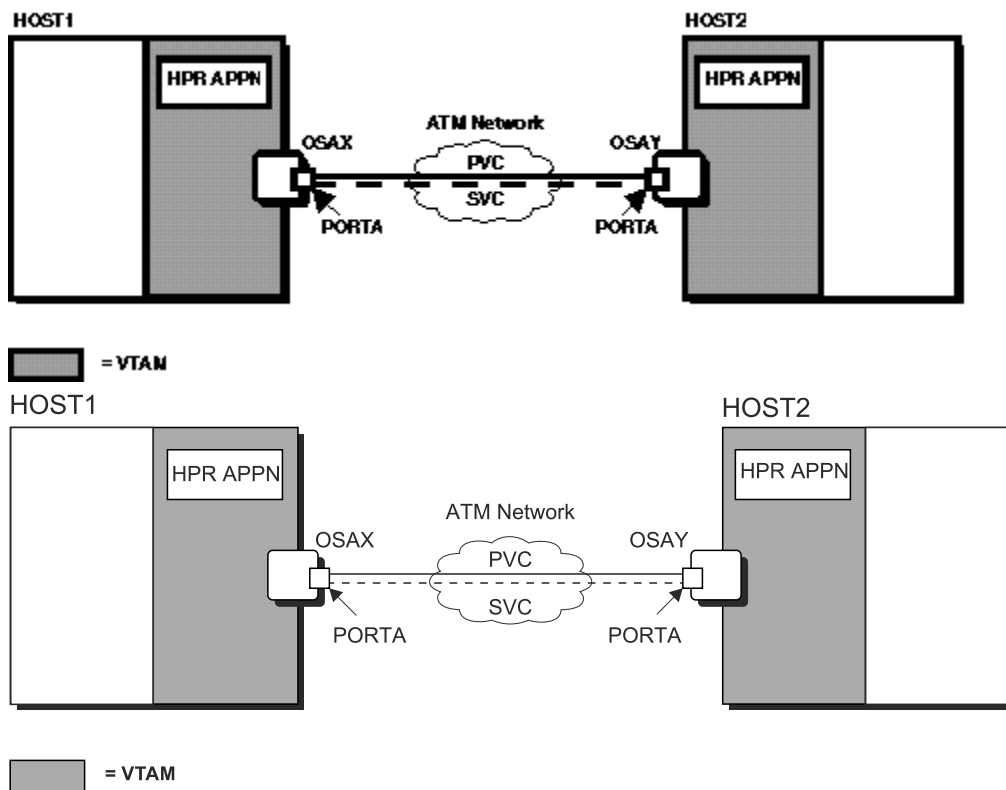


Figure 14. Basic ATM configuration

The major elements of an ATM configuration that must be defined to VTAM are:

- VTAM connection to the IBM Open Systems Adapter
- The port on the IBM Open Systems Adapter through which the ATM network is accessed
- The transmission groups (TGs) that route data:
 - Across permanent virtual channels (PVCs)
 - Across switched virtual channels (SVCs)
 - To connection networks

Defining VTAM connection to the IBM Open Systems Adapter

Figure 15 on page 56 illustrates a VTAM connection to the IBM Open Systems Adapter.




Figure 15. VTAM connection to the IBM Open Systems Adapter

VTAM communicates with the IBM Open Systems Adapter using an HPDT MPC connection. For more information about HPDT MPC connections, see [“Multipath channel connections”](#) on page 40.

Define the following key characteristics of the MPC connection in the transport resource list (TRL) major node:

- Characteristics you must specify
 - Name of the IBM Open Systems Adapter
 - Subchannel addresses of the WRITE and READ paths
 - Name of the port associated with the IBM Open Systems Adapter to which the WRITE and READ paths are connected
- Characteristics specified by default
 - Line control (LNCTL=MPC)

LNCTL=MPC is required for ATM native connections and is automatically defined for you when you code the PORTNAME operand on the TRLE definition statement in the TRL major node. The only possible value for ATM native connections is LNCTL=MPC.
 - MPC level (MPCLEVEL=HPDT)

MPCLEVEL=HPDT is required for ATM native connections and is automatically defined for you when you code the PORTNAME operand on the TRLE definition statement in the TRL major node. The only possible value for ATM native connections is MPCLEVEL=HPDT.

In addition to required user-specified characteristics, optional characteristics can be defined. See the [z/OS Communications Server: SNA Resource Definition Reference](#) for information about how to code the definition statements and operands used to define both the key and optional characteristics.

The example shown in [Figure 16 on page 57](#) is based on the configuration in [Figure 15 on page 56](#) and represents definitions in the VTAMLST data set for the VTAM in HOST1. Following [Figure 16 on page 57](#) are descriptions of the major nodes, definition statements, and operands used in [Figure 16 on page 57](#) to define VTAM connection to the IBM Open Systems Adapter.

TRL1	VBUILD	TYPE=TRL
OSAX	TRLE	WRITE=(501), READ=(500), PORTNAME=PORTA, LNCTL=MPC, MPCLEVEL=HPDT, . . .

Figure 16. Definition of VTAM connection to the IBM Open Systems Adapter

Definition statements	Description
Name of the IBM Open System Adapter	<p>The name field of the TRLE definition statement indicates that the name of the IBM Open Systems Adapter is OSAX. The name specified here is the name of the IBM Open Systems Adapter.</p> <p>The name of the IBM Open Systems Adapter is defined during IBM Open Systems Adapter configuration on the ATM Native Settings panel, which is part of the OSA Configuration pull-down option accessed through the OSA/SF OS/2 interface.</p>

Definition statements	Description
Subchannel Addresses of the WRITE and READ Paths	<p>The WRITE operand on the TRLE definition statement specifies that the subchannel address of the WRITE path is 501. The READ operand specifies that the subchannel address of the READ path is 500. The last two digits of the subchannel address of the READ path must match the Even Unit Address. The Even Unit Address and the ATM OSA-2 physical port are defined during IBM Open Systems Adapter configuration on the ATM Native Settings panel, which is part of the OSA Configuration pull-down option accessed through the OSA/SF OS/2 interface.</p> <p>Note: Only one subchannel address is specified for the READ path and one for the WRITE path. The address specified for the READ path must be an even number that is one less than the address specified for the corresponding WRITE path.</p>
Name of the Port Associated with the IBM Open Systems Adapter to which the WRITE and READ Paths are Connected	<p>The PORTNAME operand on the TRLE definition statement specifies that the name of the port associated with this IBM Open Systems Adapter—through which an ATM network can be accessed—is PORTA. The name specified here is the name of the ATM OSA-2 physical port.</p> <p>The Even Unit Address and the ATM OSA-2 physical port are defined during IBM Open Systems Adapter configuration on the ATM Native Settings panel, which is part of the OSA Configuration pull-down option accessed through the OSA/SF OS/2 interface.</p>

Defining the port on the IBM Open Systems Adapter through which the ATM network is accessed

[Figure 17 on page 59](#) shows a port example.

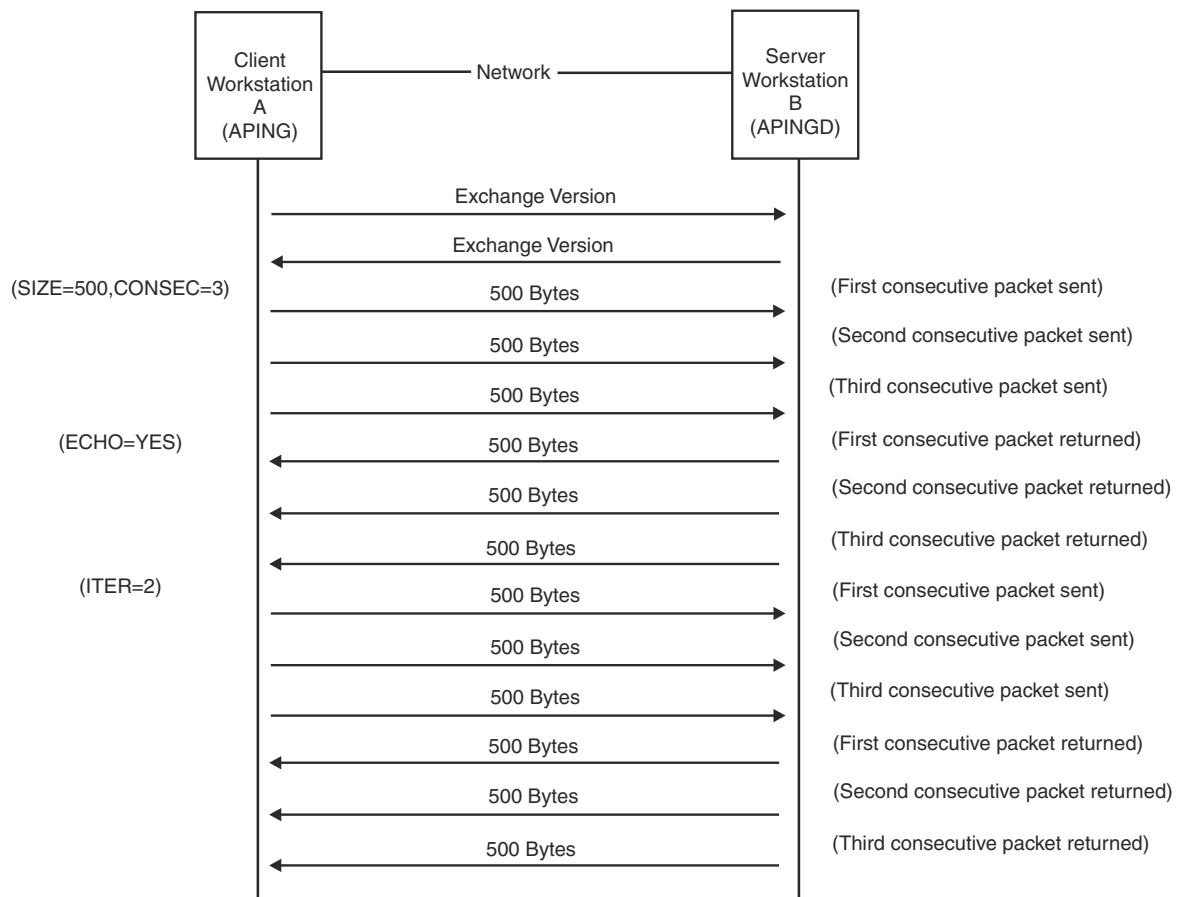


Figure 17. Port on the IBM Open Systems Adapter through which the ATM network is accessed

VTAM accesses the ATM network through a port on the IBM Open Systems Adapter. Associated with the port are links used for permanent virtual channels (PVCs) and switched virtual channels (SVCs). These PVCs and SVCs carry data across the ATM network. Define the following characteristics of the port in the external communication adapter (XCA) major node:

- Characteristic you must specify is the name of the port.
- Characteristic specified by default is the type of shared access transport facility (SATF) accessed through the port (MEDIUM=ATM).

MEDIUM=ATM is required for ATM native connections and is automatically defined for you when you do not code the ADAPNO and CUADDR operands on the PORT definition statement in the XCA major node. The only possible value for ATM native connections is MEDIUM=ATM.

See the [z/OS Communications Server: SNA Resource Definition Reference](#) for information about how to code the definition statements and operands used to define these characteristics. No other characteristics of the port need to be defined.

Figure 18 on page 59 is based on the configuration in Figure 17 on page 59 and represents definitions in the VTAMLST data set for the VTAM in HOST1. Following Figure 18 on page 59 are descriptions of the major nodes, definition statements, and operands used in Figure 18 on page 59 to define the port on the IBM Open Systems Adapter through which the ATM network is accessed.

OSAXCA1	VBUILD	TYPE=XCA
	PORT	PORTNAME=PORTA, MEDIUM=ATM

Figure 18. Definition of port on the IBM Open Systems Adapter through which the ATM network is accessed

The PORTNAME operand on the PORT definition statement specifies that the name of the port on the IBM Open Systems Adapter with which the links used for PVCs and SVCs are associated is PORTA. The name specified here must also be defined on the PORTNAME operand on a TRLE definition statement in the TRL major node. Both specifications of the port name must match the name of the ATM OSA-2 physical port.¹

Defining the APPN transmission groups that route data across the ATM network




Figure 19. TGs that route data across the ATM network

VTAM uses transmission groups (TGs) to route data across the ATM network. These TGs can be over PVCs and SVCs, as shown in [Figure 19 on page 60](#). One TG is associated with each PVC or SVC.

In networks where multiple nodes can communicate with one another across an ATM network, you can define TGs to connection networks. This can minimize the amount of definition required to establish routes among the multiple nodes.

Defining transmission groups over permanent virtual channels:

Permanent virtual channels (PVCs) represent permanent connections. They are reserved by the ATM network and are available as long as the network is active. PVCs appear to VTAM as nonswitched lines. As such, the TGs that are assigned to them are defined in groups headed by a GROUP definition statement that specifies DIAL=NO.

Because PVCs are associated with a port on the IBM Open Systems Adapter, the TGs that are assigned to them are defined in the XCA major node that defines the port with which the PVCs are associated.

Define the following key characteristics of a TG over a PVC in the XCA major node:

- Characteristic you must specify is the name of the PVC.
- Characteristics you should specify
 - Name of the remote node with which VTAM can communicate over the TG
 - Route calculation characteristics
 - Cost per connect time
 - Cost per byte
 - Security
 - Propagation delay
 - Effective capacity
 - User-defined values

Note: IBM supplies default TG profiles that define these characteristics for best effort and reserved bandwidth connections across public and private ATM networks. IBM recommends that you use these default TG profiles. For examples of the default TG profiles, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

- Characteristics specified by default

These characteristics are required for ATM native connections and are automatically defined for you when you do not code the ADAPNO and CUADDR operands on the PORT definition statement in the XCA major node, which is the equivalent of coding MEDIUM=ATM. The only possible values for ATM native connections are those shown in parentheses in the following list of characteristics.

¹ The name of the ATM OSA-2 physical port is defined during IBM Open Systems Adapter configuration on the ATM Native Settings panel, which is part of the OSA Configuration pull-down option accessed through the OSA/SF OS/2 interface.

- PU type of the remote node (PUTYPE=2)
- Type of connection (CONNTYPE=APPN)
- High-performance routing (HPR) enablement (HPR=YES)
- Channel contact procedure (XID=YES)

Figure 20 on page 61 shows the definition of a TG over a PVC. Following Figure 20 on page 61 are descriptions of the major nodes and the definition statements and operands used in Figure 20 on page 61 to define a TG over a PVC.

OSAXCA1	VBUILD	TYPE=XCA
	PORT	PORTNAME=PORTA, MEDIUM=ATM
	GROUP	DIAL=NO
PVCLN1	LINE	PVCNAME=PVC1, . . .
PVCPU1	PU	CPNAME=HOST2, TGP=ATMPVCCB, CONNTYPE=APPN, HPR=YES, PUTYPE=2, XID=YES . . .

Figure 20. Definition of a TG over a PVC

Name of the PVC: The PVCNAME operand on the LINE definition statement specifies that the name of the PVC is PVC1. The name specified here is the name of the PVC associated with the ATM OSA-2 physical port.² The ATM network provider must configure the ATM network to provide the PVC connection.

Note: The IBM Open Systems Adapter limits the total number of PVCs that can be associated with one port. See <https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-z-systems-express-customers-guide-reference> for information about this limitation.

Name of the remote node with which VTAM can communicate over the TG: The CPNAME operand on the PU definition statement specifies that the name of the remote node with which VTAM can communicate over this TG is HOST2.

Route calculation characteristics:

You can specify route calculation characteristics on the PU definition statement on the following operands:

- COSTTYPE
- COSTBYTE
- SECURITY
- PDELAY
- CAPACITY
- UPARM1
- UPARM2
- UPARM3

² The name of the PVC is defined during IBM Open Systems Adapter configuration on the ATM Native Settings panel, which is part of the OSA Configuration pull-down option accessed through the OSA/SF OS/2 interface.

Or, they can be determined by the TGP operand on the PU definition statement, which specifies the name of an IBM-supplied APPN TG profile definition. In [Figure 20 on page 61](#), they are determined by the TGP operand, which specifies the name of the IBM-supplied profile definition for ATM TGs over PVCs, campus best effort.

Defining transmission groups over switched virtual channels:

[Figure 21 on page 62](#) illustrates how TGs are defined over switched virtual channels.

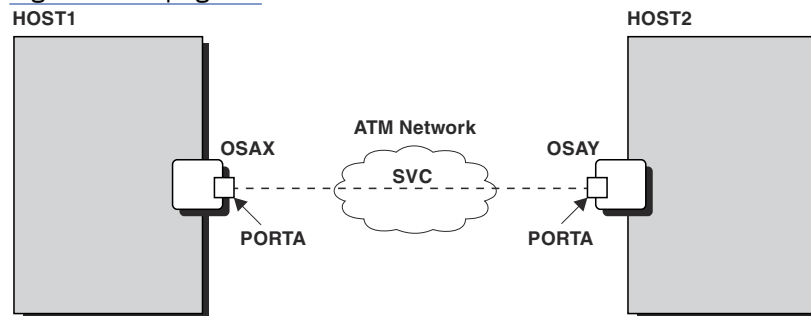


Figure 21. TG over an SVC

Switched virtual channels represent temporary connections. They are established through a dial operation and are available for as long as the connection is required. SVCs appear to VTAM as switched lines. As such, the TGs that are assigned to them are defined in groups headed by a GROUP definition statement that specifies DIAL=YES.

Because SVCs are associated with a port on the IBM Open Systems Adapter, the TGs that are assigned to them are defined in the XCA major node that defines the port with which the SVCs are associated.

Define the following key characteristics of a TG over an SVC in the XCA major node:

- Characteristic you must specify is the maximum number of simultaneous connections using the port
- Characteristics you should specify
 - Whether the line used for the TG can be used for calls initiated by a remote node, VTAM or both
 - Whether PUs are to be dynamically created during call-in operations or explicitly defined in the switched major node

If a line used for a TG over an SVC is used for calls initiated by VTAM (CALL=OUT or CALL=INOUT), define the following key characteristics of the TG in the switched major node:

- Characteristics you must specify
 - The type of APPN routes that can use the SVC
 - The name of the group that contains the dial-out line definitions for the TG
 - Name of the remote node with which VTAM can communicate over the TG
 - Address through which the remote node can be reached
 - ATM channel characteristics
 - Best effort indicator
 - Cell rates
 - Traffic management options
- Characteristics you should specify
 - Route calculation characteristics
 - Cost per connect time
 - Cost per byte
 - Security

- Propagation delay
- Effective capacity
- User-defined values

Note: IBM supplies default TG profiles that define these characteristics for best effort and reserved bandwidth connections across public and private ATM networks. It is recommended that you use these default TG profiles. For examples of the default TG profiles, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

- ATM channel characteristics including the Quality of Service (QoS) class

In addition to required characteristics, optional characteristics can be defined. See the [z/OS Communications Server: SNA Resource Definition Reference](#) for information about how to code the definition statements and operands used to define both the key and optional characteristics.

Figure 22 on page 63 is based on the configuration in Figure 21 on page 62 and represents definitions in the VTAMLST data set for the VTAM in HOST1. Following Figure 22 on page 63 are descriptions of the major nodes and the definition statements and operands used in Figure 22 on page 63 to define a TG over an SVC.

OSAXCA1	VBUILD	TYPE=XCA
	PORT	PORTNAME=PORTA, MEDIUM=ATM
SVCGRP1	GROUP	DIAL=YES, CALL=INOUT, DYNPU=YES, .
		.
SVCLN1	LINE	.
		.
SVCPU1	PU	.
		.
		.
SVCLN2	LINE	.
		.
		.
SVCPU2	PU	.
		.
		.
OSASWT1	VBUILD	TYPE=SWNET
SWTPU1	PU	CPNAME=HOST2, TGP=ATMSVCCR, .
		.
		.
	PATH	GRPNM=SVCGRP1, DLCADDR=(1,C,ATMSVC,EXCLUSIVE), DLCADDR=(7,BCD,03,00, 00000191,00, 00000191,00), DLCADDR=(8,X,03,03,03), DLCADDR=(21,X,0002, 3999999999999999), 9998010131504553543460), .
		.
		.
		.

Figure 22. Definition of a TG over an SVC

Maximum number of simultaneous connections:

The number of sets of LINE and PU definition statements in the XCA major node represents the maximum number of simultaneous connections that can be established using the port. In [Figure 22 on page 63](#), the maximum number of simultaneous connections is 2. The LINE definition statements serve as placeholders for lines used to connect to remote nodes that are defined either dynamically or in a switched major node, and are used when those nodes are activated.

Note: The IBM Open Systems Adapter limits the total number of SVCs that can be associated with one port. See <https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-z-systems-express-customers-guide-reference> for information about this limitation.

Who can initiate calls:

The CALL operand on the GROUP definition statement in the XCA major node specifies that the lines in this group can be used for calls initiated by both the VTAM in HOST1 and remote nodes.

Dynamic creation of PUs:

The DYNPU operand on the GROUP definition statement in the XCA major node specifies that PUs are to be dynamically created when remote nodes call in to VTAM. By enabling PUs to be dynamically created, you eliminate the need to code PU and PATH definition statements in the switched major node.

The type of APPN routes that can use the SVC:

In Figure 22 on page 63, EXCLUSIVE on the DLCADDR operand with a subfield of 1 on the PATH definition statement in the switched major node indicates that the TG can be assigned to an SVC that is used only for this TG, and it cannot be assigned to an SVC that is used for connection network TGs.

Meanings of the values coded on the DLCADDR operand are directly related to ATM channel characteristics. For additional information about what the values indicate, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Name of the group that contains the dial-out line definitions for the TG:

The GRPNM operand on the PATH definition statement in the switched major node specifies that the line to be used for the TG is defined in the group named SVCGRP1 in the XCA major node. This group follows a PORT definition statement, and identifies the lines to be used for the TG as those associated with the port defined on that PORT statement (PORTA in Figure 22 on page 63).

Name of the remote node with which VTAM can communicate over the TG:

The CPNAME operand on the PU definition statement specifies that the name of the remote node with which VTAM can communicate over this TG is HOST2.

Address through which the remote node can be reached:

The DLCADDR operand with subfield 21 on the PATH definition statement in the switched major node specifies that the 20-byte address through which HOST2 can be reached is identified as 399999999999999999998010131504553543460 and that this is an International Organization for Standardization (ISO) network service access point (NSAP) address, indicated by the 0002 preceding the address. The address through which HOST2 can be reached was obtained through configuration of OSAY in [Figure 22 on page 63](#). OSAY is the IBM Open Systems Adapter through which HOST2 is connected to the ATM network. When the remote node is a VTAM node, configuration of the IBM Open Systems Adapter through which the remote node is connected should be completed before definition of the calling VTAM node is attempted. See <https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-z-systems-express-customers-guide-reference> for information about how to determine the ATM address through which the remote VTAM node can be reached.

Note: You can also determine the ATM address through which the remote VTAM node can be reached by starting the remote VTAM, then activating the following resources in the remote VTAM:

- The TRL major node
- The XCA major node that defines the port on the IBM Open Systems Adapter through which the remote VTAM can be reached
- Any line defined in the XCA major node

After activating the remote VTAM resources, issue a DISPLAY ID command for the XCA major node. The resulting display shows the ATM address for the remote VTAM node.

You can follow this same procedure to obtain the ATM address for the local VTAM as well.

Route calculation characteristics:

You can specify route calculation characteristics on the PU definition statement on the following operands:

- CAPACITY
- COSTBYTE
- COSTTYPE
- PDELAY
- SECURITY
- UPARM1
- UPARM2
- UPARM3

Or, they can be determined by the TGP operand on the PU definition statement, which specifies the name of an IBM-supplied APPN TG profile definition. In [Figure 22 on page 63](#), they are determined by the TGP operand, which specifies the name of the IBM-supplied profile definition for ATM TGs over SVCs (campus reserved bandwidth).

Best effort indicator, cell rates, and traffic management options:

The DLCADDR operand with a subfield of 7 specifies the following values:

03

The format of the best effort indicator, cell rates, and traffic management options is defined by the ATM network.

00

Guaranteed bandwidth is required.

00006000

The forward peak cell rate (cells per second), in binary coded decimal, for cell loss priority 0+1 (CLP=0+1) is 6000.

00004000

The forward sustainable cell rate (cells per second), in binary coded decimal, for (CLP=0) is 4000.

00000191

The forward maximum burst size (cells), in binary coded decimal, for (CLP=0) is 191.

00

Tagging is not requested in the forward direction.

00006000

The backward peak cell rate (cells per second), in binary coded decimal, for cell loss priority 0+1 (CLP=0+1) is 6000.

00004000

The backward sustainable cell rate (cells per second), in binary coded decimal, for (CLP=0) is 4000.

00000191

The backward maximum burst size (cells), in binary coded decimal, for (CLP=0) is 191.

00

Tagging is not requested in the backward direction.

Quality of Service (QoS) class:

The DLCADDR operand with a subfield of 8 specifies the following values:

Value

Meaning

03

The format of the QoS class is defined by the ATM network.

03


The QoS class is connection-oriented for outbound (forward) data traffic.

The QoS class is connection-oriented for inbound (backward) data traffic.

Defining transmission groups to connection networks

Figure 23 on page 66 shows an ATM configuration in which multiple nodes can communicate with one another across an ATM network. HOST2 and HOST3 are end nodes that dial in to the network node server, HOST1. All nodes can establish switched connections with one another.

To enable optimal route selection among these nodes, an extensive number of TGs needs to be defined. Specifically, each node needs to define TGs to every other node. So, in Figure 23 on page 66, HOST1 needs to define TGs to HOST2 and HOST3, HOST2 needs to define TGs to HOST1 and HOST3, and HOST3 needs to define TGs to HOST1 and HOST2. In a large configuration with many nodes connected by many SVCs, the definition of TGs can become overwhelming.



In this example, DELAY is equal to t . This time period is large enough to allow four messages to arrive at the NCP. The time t has expired, and the NCP presents an attention interrupt. VTAM reads all four messages, having received only one attention interrupt.

Figure 23. Multiple nodes communicating across an ATM network

APPN's connection network function reduces this extensive TG definition for connections among multiple nodes across an ATM network. A connection network is a representation of a shared access transport facility, such as an ATM network, that handles the routing of data among the nodes communicating across the shared access transport facility. It does this by enabling the shared access transport facility to be defined as a virtual node. As a result, end nodes need to define TGs only to the virtual node and to the network node server. The network node server, though it does not always need to, should also define a TG to the virtual node.

Figure 24 on page 67 shows the ATM configuration in Figure 23 on page 66 with a virtual node. The virtual node represents the ATM network to the nodes that can communicate with one another across the ATM network.




Figure 24. ATM configuration with a connection network

Now, HOST1 needs to define a TG only to VNODE1, as does HOST2 and HOST3. HOST2 and HOST3 must also each define a TG to HOST1 because HOST1 is their network node server. The same routing capability achieved by six TG definitions for the configuration in Figure 23 on page 66 is now achieved by five TG definitions using the virtual node in Figure 24 on page 67. In large networks with many nodes connected by many SVCs, the reduction in TG definitions is much more significant. For example, without the use of a connection network, a configuration with five end nodes that dial into one network node server requires 30 TG definitions to enable optimal route selection. With the use of a connection network, that same configuration requires only 11 TG definitions to enable the same optimal route selection.

If another network node server and its associated end nodes wanted to communicate with HOST1, HOST2, and HOST3 through the same connection network, each new end node would need to define a TG to the virtual node and to the new network node server. The new network node server would also need to define a TG to the virtual node and to HOST1.

Multiple TGs through one port can be to the same virtual node. In addition, TGs to multiple virtual nodes can be through the same port.

Connections through connection networks occur over SVCs, which appear to VTAM as switched lines. Thus, TGs to connection networks are defined in groups headed by a GROUP definition statement that specifies DIAL=YES.

Because SVCs are associated with a port on the IBM Open Systems Adapter, TGs to connection networks are defined in the XCA major node that defines the port with which the SVCs are associated.

To define the following key characteristics of a TG to a connection network in the XCA major node:

- Characteristics you must specify
 - The type of APPN routes that can use the SVC
 - Name of the virtual node that represents the ATM network
 - The maximum number of simultaneous connections using the port
 - ATM channel characteristics
 - Best effort indicator
 - Cell rates
 - Traffic management options
- Characteristics you should specify
 - Whether the line used for the TG can be used for calls initiated by a remote node, VTAM, or both

- Whether PUs are to be dynamically created during call-in operations or explicitly defined in the switched major node
- Route calculation characteristics
 - Cost per connect time
 - Cost per byte
 - Security
 - Propagation delay
 - Effective capacity
 - User-defined values

Note: IBM supplies default TG profiles that define these characteristics for best effort and reserved connections across public and private ATM networks and recommends use of these default TG profiles. For examples of the default TG profiles, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

- ATM channel characteristics
 - Quality of Service (QoS) class
- Characteristics specified by default
 - High-performance routing (HPR) enablement (HPR=YES)

HPR=YES is required for ATM native connections and is automatically defined for you when you do not code the ADAPNO and CUADDR operands on the PORT definition statement in the XCA major node, which is the equivalent of coding MEDIUM=ATM. The only possible value for ATM native connections is HPR=YES.

In addition to the required characteristics, other optional characteristics can be defined. See the [z/OS Communications Server: SNA Resource Definition Reference](#) for information about how to code the definition statements and operands used to define both the key and optional characteristics.

[Figure 25 on page 69](#), [Figure 26 on page 70](#), and [Figure 27 on page 71](#) are based on the configuration in [Figure 24 on page 67](#) and represent definitions in the VTAMLST data sets for the VTAMs in HOST1, HOST2, and HOST3. Following [Figure 27 on page 71](#) are descriptions of the major nodes, definition statements, and operands used in [Figure 25 on page 69](#), [Figure 26 on page 70](#), and [Figure 27 on page 71](#) to define TGs to a connection network.

```

HOST1

          OSAXCA1  VBUILD  TYPE=XCA
                      PORT  PORTNAME=PORTA,
                          MEDIUM=ATM

          CNGRP1   GROUP   DIAL=YES,
                          CALL=INOUT,
                          DLCADDR=(1,C,ATMSVC, &cont
VNODE1,EXCLUSIVE),
00006000,
00),
                                DLCADDR=(7,BCD,03,00, &cont
                                00004000,00000191, &cont
                                DLCADDR=(8,X,03,03),
                                DYNPU=YES,
                                TGP=ATMSVCCR,
                                HPR=YES,
                                .
                                .
                                .

          CNLN1    LINE
          CNPU1    PU
          CNLN2    LINE
          CNPU2    PU
          CNLN3    LINE
          CNPU3    PU

```

Figure 25. Definitions in VTAMLST for the VTAM in HOST1

```

HOST2

          OSAXCA1  VBUILD  TYPE=XCA
                      PORT  PORTNAME=PORTA,
                          MEDIUM=ATM

          CNGRP1  GROUP   DIAL=YES,
                          CALL=INOUT,
                          DLCADDR=(1,C,ATMSVC,
VNODE1,EXCLUSIVE),
00006000,
00),
                          DLCADDR=(7,BCD,03,00,
                                00004000,00000191,
                          DLCADDR=(8,X,03,03),
                          DYNPU=YES,
                          TGP=ATMSVCCR,
                          HPR=YES,
                          .
                          .
                          .

          CNLN1   LINE
          CNPU1   PU
          CNLN2   LINE
          CNPU2   PU
          CNLN3   LINE
          CNPU3   PU

          OSASWT1  VBUILD  TYPE=SWNET
          SWTPU1   PU       CPNAME=HOST1,
                          TGP=ATMSVCCR,
                          .
                          .
                          .

                          PATH  GRPNM=CNGRP1,
                          DLCADDR=(1,C,ATMSVC,EXCLUSIVE),
                          DLCADDR=(7,BCD,03,00,
00006000,00004000,
                                00000191,00,
00006000,00004000,
                                00000191,00),
                          DLCADDR=(8,X,03,03,03),
                          DLCADDR=(21,X,0002,
                                399999999999999999
9998010131504553543460),
                          .
                          .
                          .

```

Figure 26. Definitions in VTAMLST for the VTAM in HOST2

```

HOST3

          OSAXCA1  VBUILD  TYPE=XCA
                      PORT  PORTNAME=PORTA,
                          MEDIUM=ATM

          CNGRP1  GROUP  DIAL=YES,
                          CALL=INOUT,
                          DLCADDR=(1,C,ATMSVC,
VNODE1,EXCLUSIVE),
                          DLCADDR=(7,BCD,03,00,
00006000,
                          00004000,00000191,
00),
                          DLCADDR=(8,X,03,03),
                          DYNPU=YES,
                          TGP=ATMSVCCR,
                          HPR=YES,
                          .
                          .
                          .

          CNLN1   LINE
          CNPU1   PU
          CNLN2   LINE
          CNPU2   PU
          CNLN3   LINE
          CNPU3   PU
          OSASWT1 VBUILD  TYPE=SWNET
          SWTPU1  PU      CPNAME=HOST1,
                          TGP=ATMSVCCR,
                          .
                          .
                          .

                          PATH  GRPNM=CNGRP1,
                              DLCADDR=(1,C,ATMSVC,EXCLUSIVE),
                              DLCADDR=(7,BCD,03,00,
00006000,00004000,
                              00000191,00,
00006000,00004000,
                              00000191,00),
                              DLCADDR=(8,X,03,03,03),
                              DLCADDR=(21,X,0002,
                              3999999999999999
9998010131504553543460),
                              .
                              .
                              .

```

Figure 27. Definitions in VTAMLST for the VTAM in HOST3

The type of APPN routes that can use the SVC:

In Figure 25 on page 69 through Figure 27 on page 71, EXCLUSIVE on the DLCADDR operand with a subfield of 1 on the PATH definition statement in the XCA major node indicates that separate TGs and SVCs can be established. These handle sessions over connection-network connections only. Sessions through connection networks cannot be routed over an SVC already established to the same remote node if that SVC was not established through a connection network.

Meanings of the values coded on the DLCADDR operand are directly related to ATM channel characteristics. For information about what the values indicate, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Name of the virtual node that represents the ATM network:

The DLCADDR operand, with subfield 1 on the GROUP definition statement in the XCA major node, specifies the name of the virtual node representing the ATM network to VTAM is VNODE1.

Maximum number of simultaneous connections:

The number of sets of LINE and PU definition statements in the XCA major node represents the maximum number of simultaneous connections that can be established using the port. In Figure 22 on page 63, the maximum number of simultaneous connections is two. The LINE definition

statements serve as placeholders for lines used to connect to remote nodes that are defined either dynamically, or in a switched major node, and are used when those nodes are activated.

Note: The IBM Open Systems Adapter limits the total number of SVCs that can be associated with one port. See <https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-z-systems-express-customers-guide-reference> for information about this limitation.

Who can initiate calls:

The CALL operand on the GROUP definition statement in the XCA major node specifies that the lines in this group can be used for calls initiated by both the VTAM in HOST1 and remote nodes.

Dynamic creation of PUs:

The DYNPU operand on the GROUP definition statement in the XCA major node specifies that PUs are to be dynamically created when remote nodes call in to VTAM. By enabling PUs to be dynamically created, you eliminate the need to code PU and PATH definition statements in the switched major node. If you explicitly define PUs in the switched major node, you lose the major benefit of the connection network: reduction of resource definition.

Route calculation characteristics:

Route calculation characteristics can be specified on the GROUP definition statement in the XCA major node on the following operands:

- CAPACITY
- COSTBYTE
- COSTTYPE
- PDELAY
- SECURITY
- UPARM1
- UPARM2
- UPARM3

They can also be determined by the TGP operand on the GROUP definition statement, which specifies the name of an IBM-supplied APPN TG profile definition. In [Figure 25 on page 69](#) through [Figure 27 on page 71](#), the route calculation characteristics are determined by the TGP operand, which specifies the name of the IBM-supplied profile definition for ATM TGs over SVCs (campus reserved bandwidth).

Best effort indicator, cell rates, and traffic management options:

The DLCADDR operand with a subfield of 7 specifies the following values:

Value

Meaning

03

The format of the best effort indicator, cell rates, and traffic management options is defined by the ATM network.

00

Guaranteed bandwidth is required.

00006000

The forward peak cell rate (cells per second) for cell loss priority 0+1 (CLP=0+1) is 6000.

00004000

The forward sustainable cell rate (cells per second) for (CLP=0) is 4000.

00000191

The forward maximum burst size (cells) for (CLP=0) is 191.

00

Tagging is not requested in the forward direction.

Quality of Service (QoS):

The DLCADDR operand with a subfield of 8 specifies the following values:

Value	Meaning
-------	---------

03	The format of the QoS class is defined by the ATM network.
03	The format of the QoS class is defined by the ATM network.
03	The QoS class is connection-oriented for outbound (forward) data traffic.

Establishing a TG over a permanent virtual channel (PVC)

For a PVC connection to be established, both nodes must activate the PVC. To activate the PVC from VTAM, follow these steps:

1. Accept the default value or specify YES on the DYNADJCP start option in the VTAM start list. Or, activate the adjacent control point (ADJCP) major node that defines the remote node with which VTAM can communicate over the TG assigned to this PVC.

```
ATCSTR01 DYNADJCP=YES
```

or

```
VARY ACT ID=ADJCP1
```

2. Activate the transport resource list (TRL) major node that defines VTAM connection to the IBM Open Systems Adapter.

```
VARY ACT ID=TRL1
```

3. Activate the external communication adapter (XCA) major node that defines the port on the IBM Open Systems Adapter used to access the ATM network.

```
VARY ACT ID=OSAXCA1
```

4. Activate the LINE and PU definition statements in the XCA major node that define the line to which the PVC connection can be assigned and the remote node with which VTAM can communicate over the PVC connection.

```
VARY ACT ID=PVCLN1
```

```
VARY ACT ID=PVCPU1
```

Establishing a TG over a switched virtual channel (SVC)

For an SVC connection to be established, both nodes must be prepared to call or receive calls.

To prepare VTAM to call or receive calls, follow these steps:

1. Accept the default value or specify YES on the DYNADJCP start option in the VTAM start list. Or, activate the ADJCP major node that defines the remote node with which VTAM can communicate over the TG assigned to this SVC.

```
ATCSTR01 DYNADJCP=YES
```

or

```
VARY ACT ID=ADJCP1
```

2. Activate the TRL major node that defines VTAM connection to the IBM Open Systems Adapter.

```
VARY ACT ID=TRL1
```

3. Activate the XCA major node that defines the port on the IBM Open Systems Adapter used to access the ATM network.

```
VARY ACT ID=OSAXCA1
```

4. Activate the LINE definition statements in the XCA major node that serve as placeholders for the lines used to connect to the remote nodes with which VTAM can communicate over the SVC connection.

```
VARY ACT ID=SVCLN1
```

```
VARY ACT ID=SVCLN2
```

The LINE definition statements in the XCA major node serve as placeholders for lines used to connect to remote nodes that are defined either dynamically, or in a switched major node, and are used when those nodes are activated.

To call a remote node from VTAM, follow these steps:

1. Activate the switched major node that defines the remote node with which VTAM can communicate over the TG assigned to the SVC.

```
VARY ACT ID=OSASWT1
```

2. Activate the physical units defined in the switched major node.

```
VARY ACT ID=SWTPU1
```

3. Dial the physical units defined in the switched major node.

```
VARY DIAL ID=SWTPU1
```

Establishing a TG to a connection network

To establish a connection-network connection, follow these steps:

1. Accept the default value or specify YES on the DYNADJCP start option in the VTAM start list. Or, activate the ADJCP major node that defines the remote nodes with which VTAM can communicate through the connection network.

```
ATCSTR01 DYNADJCP=YES
```

or

```
VARY ACT ID=ADJCP1
```

2. Activate the TRL major node that defines VTAM connection to the IBM Open Systems Adapter.

```
VARY ACT ID=TRL1
```

3. Activate the XCA major node that defines the port on the IBM Open Systems Adapter used to access the ATM network.

```
VARY ACT ID=OSAXCA1
```

4. Activate the LINE definition statements in the XCA major node that serve as placeholders for the SVCs to the remote nodes with which VTAM can communicate over the connection-network connection.

```
VARY ACT ID=CNLN1
```

```
VARY ACT ID=CNLN2
```

```
VARY ACT ID=CNLN3
```


The LINE definition statements in the XCA major node serve as placeholders for SVCs to remote nodes that are defined either dynamically, or in a switched major node, and are used when sessions to those nodes are established.

APPN multiple network connectivity

APPN multiple network connectivity support enables connectivity between APPN networks, without the exchange of topology information between the networks. APPN multiple network connectivity support may exist between APPN networks that have different network IDs, and also between APPN subnetworks that have the same network ID. APPN multiple network connectivity support can also be used with High-Performance Routing (HPR) support to enable HPR routes between nodes in different subnetworks. See [“High-Performance Routing \(HPR\)” on page 388](#) for information about HPR support.

APPN multiple network connectivity support uses border nodes at the boundaries between APPN networks or subnetworks. With APPN multiple network connectivity support, full APPN directory and session initiation support is available across these boundaries. APPN directory search activity can be managed through several user-definable parameters for border nodes. Class of Service definitions do not need to be consistent across network or subnetwork boundaries, allowing for independent management.

APPN multiple network connectivity support enables you to eliminate topology exchange between APPN networks or subnetworks when the exchange of such data is undesirable (for example, for networks of different enterprises). Topology exchange is eliminated between border nodes. This way, you can protect smaller subnetworks from the volume of topology in a larger subnetwork or from processing searches for resources that do not exist in the smaller subnetwork. Thus, APPN multiple network connectivity support can be used to manage the increase in topology and directory search activity as an APPN network grows.

For example, the APPN network in [Figure 28 on page 76](#) is divided into four subnetworks. With APPN multiple network connectivity support, full APPN directory and session initiation support is available across all subnetwork boundaries, and yet, the creation of smaller subnetworks protects each from the volume of topology and search processing that would exist in the entire APPN network (which could, of course, extend beyond what is shown).

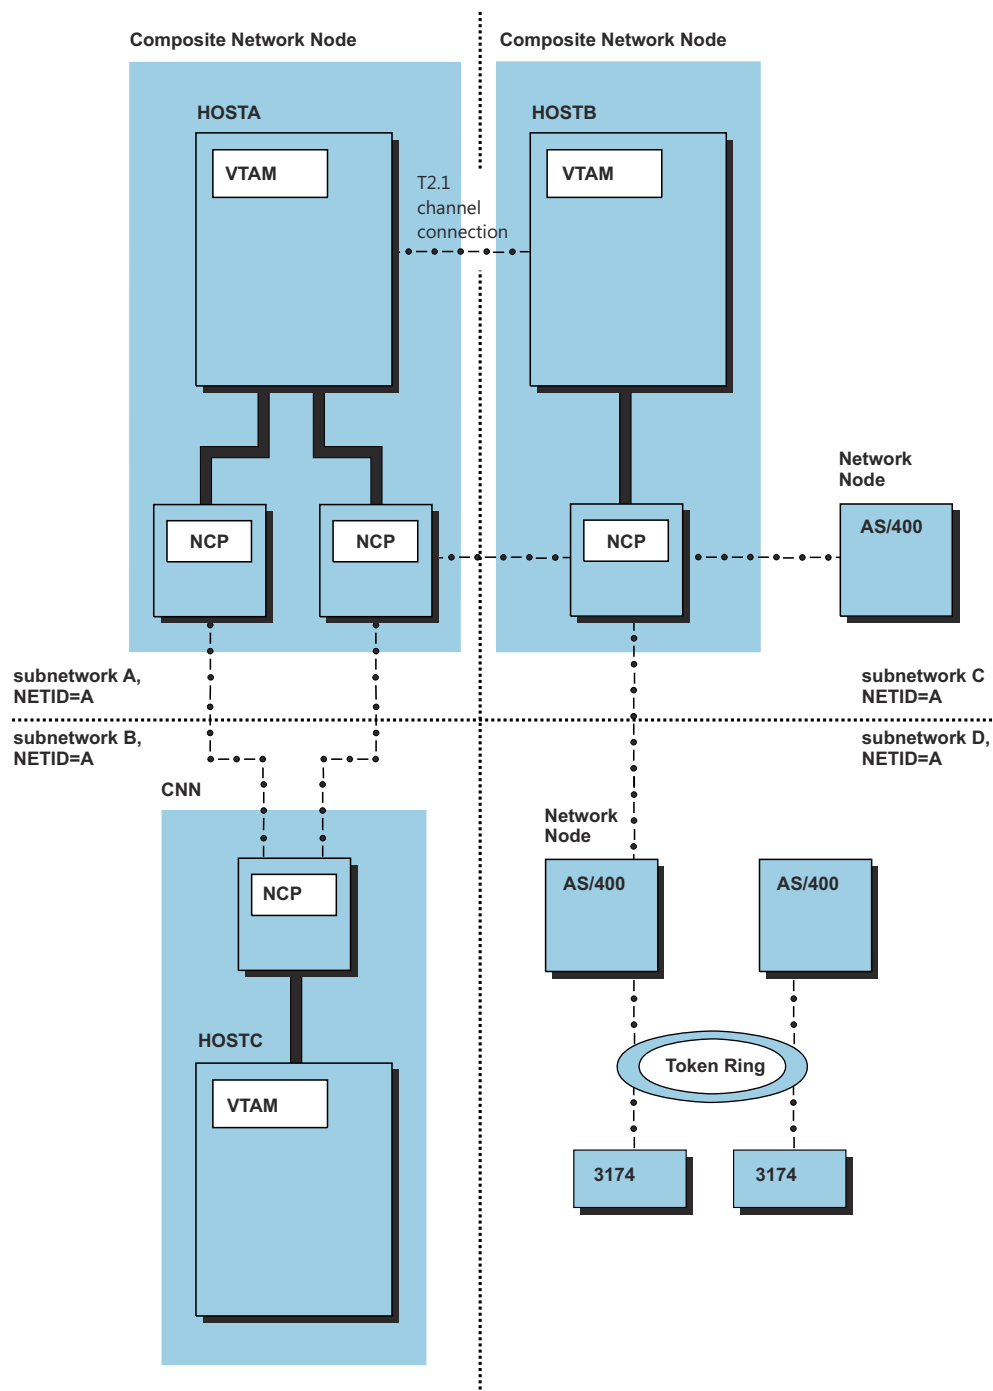


Figure 28. APPN subnetworks through APPN multiple network connectivity support

There are two types of border node boundaries:

- Peripheral subnetwork boundaries
- Extended subnetwork boundaries

Restriction: APPN multiple network connectivity support cannot be used to connect two APPN networks or subnetworks through a connection network unless the connection network uses Enterprise Extender and is defined as VNTYPE=GLOBAL. For more information, see [“Configuring the EE connection network”](#) on page 116.

Peripheral subnetwork boundaries

A peripheral subnetwork boundary is a connection over a subnetwork boundary between a border node (VTAM with APPN multiple network connectivity support enabled) and a network node with no border node function. The VTAM border node portrays an end node image to the network node, thus preventing topology information from being sent by the network node to VTAM.

VTAM offers the peripheral subnetwork boundary to provide APPN subnetwork interconnection with nodes that do not have border node function. Because the nonborder node in a peripheral boundary is not aware that the subnetwork boundary exists, compared to an extended boundary, there is considerable loss of control over search order, redundant searching, and looping search prevention. For these reasons, it is recommended that you use an extended boundary whenever possible. When using a peripheral boundary, it is wise to use a single subnetwork boundary connection to the peripheral subnetwork (the subnetwork of the nonborder node) to prevent locate requests from exiting the subnetwork over one boundary and then reentering over another boundary. If a peripheral subnetwork has more than one boundary, the user should customize routing by a user exit or routing definitions to ensure that the requests do not reenter the same peripheral subnetwork. This precaution is necessary to avoid possible failed sessions because of a looping locate path.

Guideline: Because these controls are not present for peripheral subnet boundaries, to avoid looping problems, limit connections to a single connection between networks when peripheral subnet boundaries are used.

Notes:

- A peripheral subnetwork boundary can be used as either the first or last boundary (or both) on a multiple subnetwork search path. If there is an intermediate subnetwork on a multiple subnetwork search path (that is, neither the search originator nor the search target are members of that subnetwork), only extended border nodes can be used to manage the intermediate subnetwork's boundaries.
- Not all network nodes are capable of attaching to a VTAM border node through a peripheral subnetwork boundary. Ensure the product and release support interoperability with peripheral border nodes (APPN option set 1013) before attempting to establish subnetwork boundaries between VTAM and a network node.
- It is possible, although strongly discouraged, to get a peripheral boundary between two VTAMs that both support extended border node by setting BN=NO in one of the two nodes. This is strongly discouraged because it bypasses the valuable looping and redundant search controls that are present if the connection is an extended subnetwork boundary.

Extended subnetwork boundaries

An extended subnetwork boundary is a connection between two border nodes that support extended subnetwork boundaries. Although topology flows are still restricted across the boundary, both nodes recognize that a network node exists on the other side, allowing for additional subnetwork hops and increased control over routing.

Note: Some border nodes do not support extended subnetwork boundaries. If one of two communicating border nodes does not support extended subnetwork boundaries, a peripheral subnetwork boundary is used.

APPN multiple network connectivity support

To enable APPN multiple network connectivity support, use the BN=YES start option at a VTAM network node (that is, NODETYPE=NN must also be specified). BN=YES indicates that VTAM is capable of establishing subnetwork boundaries.

Note: NCP Version 7 Release 1 or higher is required for both extended and peripheral subnetwork boundaries through an NCP.

You need to code only the BN=YES start option to take advantage of APPN multiple network connectivity support. Some controls are available for customization:

- An adjacent cluster routing definition list (start with the VBUILD TYPE=ADJCLUST definition statement) can be used to customize routing between subnetworks. The ADJCLUST definition is used when a border node determines it is unable to satisfy a search request [using information in its DS database, TRS database, or symbolic resolution table (SRT)] for a resource that is not in its domain. The border node builds a subnetwork routing list, used to control searching both the native subnetwork and nonnative subnetworks. The subnetwork routing list is built from the ADJCLUST definition, the network topology database (which is aware of all native border nodes), SNVC start option, and learned information from prior searches. The BNDYN and BNORD start options control how this information is used to build the subnetwork routing list.

The subnetwork routing list is built differently at exit border nodes (a border node receiving a request from another node in its native subnetwork) when compared to entry border nodes (a border node receiving a request across a subnetwork boundary from a nonnative node) and origin border nodes (a border node that is the OLU, or a CDS where the OLU is not a border node). An exit border node includes only adjacent nonnative network nodes in its subnetwork routing list. Entry and origin border nodes include these also, and its own name (used to search the native APPN subnetwork) and all other border nodes in the native APPN subnetwork. Any nodes in the subnetwork routing list that cannot be routed to are automatically pruned from the list.

When the subnetwork routing list is built, the border node sequentially sends a directed Locate to each node on the list, progressing down the list only if a negative reply is received. If a border node encounters its name in the subnetwork routing list, it performs searching of its locally attached subarea network and native APPN subnetwork, possibly including a network broadcast search and interchange node search. The interchange node search allows any subarea network attached to the native APPN subnetwork to be searched as well.

- The BNDYN start option is used to define how VTAM adds entries to the adjacent cluster routing table. By default, VTAM adds a limited set of nodes to the table.

Notes:

1. When coding BNDYN=NONE, you will need to code an adjacent cluster table, otherwise session requests might not leave this node. Be sure to code this node's name in the table to allow the native subnet to be searched.
 2. In order to search its own network, when BNDYN=NONE and VTAM is the NNS(OLU) or CP(OLU), a minimum of two adjacent cluster tables are needed (one to use when the NETID of the resource is known, and a default adjacent cluster table to use when the NETID of the resource is not known). Adjacent cluster tables for other NETIDs might also be needed when VTAM directory information does not reflect the actual location of the DLU. [z/OS Communications Server: SNA Resource Definition Reference](#) provides [adjacent cluster routing definitions](#).
- The BNORD start option is used to control the search order when searching across subnetwork boundaries. By default, VTAM gives preference to nodes for which the most recent search was successful and nodes whose network ID matches that of the destination logical unit (DLU). Specify BNORD=DEFINED to use the defined search order.
 - The SNVC start option is a number in the range of 1–255 that specifies the maximum number of subnetworks to which a request can be forwarded, including this subnetwork. For example, the default for the SNVC start option is 3, indicating that a request can be routed from this subnetwork across up to two subnetwork boundaries, for a total of three subnetworks visited when you count the original subnetwork. Given the connectivity shown in [Figure 28 on page 76](#), the default would not allow the routing of a request from subnetwork D to subnetwork B (SNVC=4 is needed).

SNVC processing is summarized as follows:

- When a request is received across a subnetwork boundary, a border node decrements the SNVC, if present. If the SNVC is 0 after being decremented, the request is rejected without further processing. When a request is received from another node in the same subnetwork as the border node, the border node does not decrement the SNVC.
- Before sending a request to another node, a border node compares the SNVC on the request received (after decrementing, if received over a subnetwork boundary) to the SNVC defined at the border node

for the destination node. The SNVC on the request sent is set to the lower of these two values. If no SNVC was received on the request, the defined value is used.

The SNVC defined at a border node is determined from the SNVC start option and the SNVC values specified in the adjacent cluster routing definition list, with the adjacent cluster routing definition list taking priority.

- Class of Service mapping definitions can be placed in a COS mapping table to map:
 - Nonnative COS definitions to native COS definitions
 - Native COS definitions to nonnative COS definitions
 - The APPNCOS start option is used to specify the default COS used for requests received over a subnetwork boundary with an unrecognizable COS listed. The value for APPNCOS is used only after COS mapping is attempted. The default value is NONE.
 - The CACHETI start option defines the number of minutes that routing information about a previous locate search is stored, if applicable. The default is eight minutes.
 - The ALIASRCH operand is specified on the ADJCP minor node definition for an adjacent non-native CP. The ALIASRCH operand is valid only when BN=YES is specified. Code ALIASRCH=NO if you want this border node to halt inbound ALIAS searches from the adjacent non-native node.
 - The AUTHNETS operand is specified on the ADJCP minor node definition for an adjacent non-native CP. The AUTHNETS operand is valid only when BN=YES is specified. Specify AUTHNETS= if you want to limit searches through this border node from the adjacent non-native node, based on the NETID value of the DLU resource.
 - The NATIVE operand is specified on the ADJCP definition statement in the adjacent CP major node, or on a PU definition statement defining an APPN connection. The NATIVE operand is valid only when BN=YES is specified. Code NATIVE=NO to define a subnetwork boundary between this node and the named adjacent CP, or between this node and the CP represented by the PU statement. Use NATIVE=NO when both nodes have the same network ID, but a subnetwork boundary is required. For example, in [Figure 28 on page 76](#), NATIVE=NO was used to create the subnetwork boundary between subnetwork A and subnetwork C. The NATIVE operand is required on only one side of a network or subnetwork boundary, though it is allowed on both sides as long as they do not conflict.
- Note:** Make sure to use the NATIVE operand consistently in your network. For example, do not specify a native connection (NATIVE=YES) from one node to a second node when the second node has a native connection to a third node that is nonnative (NATIVE=NO) to the first node. This is not valid.
- The RTPONLY operand is specified on the ADJCP definition statement in the adjacent CP major node. The RTPONLY operand is valid only when BN=YES is specified. Code RTPONLY=YES on the ADJCP definitions for nonnative nodes if you want this border node to maintain awareness of all sessions established to, from, or through the adjacent nonnative node being defined. RTPONLY=YES can be used only when the HPR start option specified RTP as the first operand. Because of the potential increase in storage and CPU use, RTPONLY=YES is recommended only when there is a crucial need to maintain session awareness (such as for accounting purposes).

By coding RTPONLY=YES, you are instructing VTAM to disallow use of the high performance routing (HPR) automatic network routing (ANR) function for any new RTPs that are established or path switched to, from, or through the adjacent nonnative node being defined. (Allowing the use of ANR could result in RTPs being established through this border node, thereby preventing the border node from maintaining awareness of any sessions that use these RTPs.) Instead of allowing RTPs to be established through this border node, coding RTPONLY=YES will result in VTAM forcing these RTPs to terminate on this border node or forcing the use of intermediate session routing (ISR) instead of HPR (or a combination of the two).

Restrictions:

1. Use of RTPONLY=YES at any APPN subnetwork boundary on a session setup path prevents the use of Global VRNs (GVRNs) for intersubnetwork connectivity, because using GVRNs could result in sessions being established across this subnetwork boundary without this border node maintaining awareness of these sessions.

2. Use of RTPONLY=YES can result in an increase in network traffic in the form of additional Route_Setup flows used for RTP establishment. These additional Route_Setup flows will occur only during the establishment of sessions that cross subnetwork boundaries defined with RTPONLY=YES.
3. Use of RTPONLY=YES can result in an increase in storage and CPU utilization because of VTAM maintaining awareness of these sessions and performing ISR instead of HPR/ANR for these sessions.

For more information about extended border node, including links to presentations and examples, see [the technote about extended border node](#).

Virtual-route-based transmission groups

A virtual-route-based transmission group (VR-based TG) connects two APPN-capable VTAM nodes through a subarea network and allows them to act as APPN peers. The VR-based TG (whose TG number is always 255) then functions as any other TG in an APPN network. APPN CP-CP sessions can be established between these two nodes, which allows all APPN functions to be performed through the subarea network. For example, if the two nodes are attached to disjoint APPN subnetworks, the topologies of these two subnetworks are exchanged over the CP-CP sessions, thereby combining them into one large APPN network. Resource location and session establishment can also be performed through the subarea network using APPN protocols, rather than transforming APPN flows to subarea flows and back to APPN flows. In addition, VR-based TG support can be used in conjunction with HPR support, enabling an HPR route to traverse a VR-based TG. See [“High-Performance Routing \(HPR\)” on page 388](#) for information about HPR support.

Only one VR-based TG can be defined between any two VTAM nodes, but this does not limit the use of this TG to only one virtual route (VR). In fact, the VR-based TG is used to represent all possible VRs between any two subareas in the domains of the two VTAMs. For example, in [Figure 29 on page 81](#), if an application on HOSTA starts a session with an application on HOSTB, the VR-based TG can be selected as the route for this session. When this occurs, VR1 can be assigned as the underlying VR for the session, which will use the channel-to-channel connection between the two hosts. If an independent LU on network node A starts a session with an independent LU on network node D, the VR-based TG can be selected for a portion of the session route. However, in this case, VR2 can be assigned as the underlying VR, so that the direct connection between the two NCPs will be used.




Figure 29. VR-based TG between composite network nodes

Because an underlying VR is always assigned to sessions established using VR-based TGs, both APPN and subarea functions must be enabled on both nodes. This means that each node must be defined as an interchange node or a migration data host. (That is, the NODETYPE start option must be coded and each host must define a unique HOSTSA number.) Also, because a limited number of subarea flows must be sent between these two nodes to determine which underlying VR should be used, VR-based TG support also depends on a CDRM session being active between the two nodes.

Notes:

1. It is possible for a session route to be calculated that includes two or more consecutive VR-based TGs, as shown in Figure 30 on page 82. When this occurs, consecutive VR-based TGs are merged together into one VR-based TG. This is referred to as RSCV pruning. The resulting VR-based TG appears to directly connect the first and last VTAM nodes (the endpoints of the consecutive VR-based TGs), even though a VR-based TG was not actually defined between these two nodes. For this reason, it is not sufficient to have active CDRM sessions only between VR-based TG partner nodes. The existing subarea requirement that every SSCP in the network have an active CDRM session with every other SSCP in that network still holds.
2. VR-based TGs cannot be used across APPN or subarea network boundaries (that is, between nodes with different network IDs) or across APPN subnetwork boundaries created using the NATIVE=NO operand. For information about APPN subnetwork boundaries created using the NATIVE=NO operand, see [“APPN multiple network connectivity”](#) on page 75.




Figure 30. Multiple contiguous VR-based TGs

Defining a VR-based TG

To define a VR-based TG between two VTAM APPN nodes, code `VRTG=YES` as a start option at both VTAMs or on the CDRM definition statement for the adjacent VTAM. If `VRTG=YES` is coded at both VTAMs, then a VR-based TG is activated automatically when the CDRM session with the adjacent VTAM is established. If the VR-based TG connects two VTAM network nodes, a topology database update (TDU) is broadcast throughout the network, so that the VR-based TG is added to the topology database of all network nodes. If either VTAM is defined as an end node, then the VR-based TG is registered with the end node network node server. Similarly, when the CDRM session between these two VTAMs terminates, the VR-based TG is deactivated and the appropriate nodes are notified that the TG is no longer active (either by sending another TDU through the network or by deleting a registered TG with an end node network node server).

Note: To enable VR-based TGs with a specific adjacent CDRM without disrupting existing LU-LU sessions, use the `SAVESESS` operand when deactivating the adjacent CDRM. Keep in mind, however, that deactivation of a CDRM with `SAVESESS` disassociates the active sessions from the CDRM. Subsequent activations and deactivations of the CDRM have no effect on these sessions.

If there are no CP-CP sessions active between the two VTAM nodes, CP-CP session establishment is automatically initiated when a VR-based TG is activated. The CP-CP sessions might use the VR-based TG or any other APPN link between the two VTAMs that supports CP-CP sessions. If CP-CP sessions are not required over a VR-based TG, code `VRTGCPCP=NO` as a start option or on the CDRM definition statement for the adjacent VTAM. Note, however, that an alternate CP-CP session path must exist between the two VTAMs, so that network traffic can be sent. For example, in [Figure 29 on page 81](#), if a connection existed between network node B and network node C and CP-CP sessions were established between them, `VRTGCPCP=NO` could be used to prevent CP-CP sessions from being established over the VR-based TG between the two VTAMs. When a session between two VTAM applications is initiated, the network traffic sent to locate the target resource would have to be sent through network node B and network node C, but the session could still be established over the direct VR-based TG connection.

The choice of which APPN TGs can be used for a given session is based on the APPN Class of Service and the characteristics defined for the links in the APPN network. As such, APPN link characteristics must be defined for all VR-based TGs as well. This is done by coding the `TGP` operand on the adjacent CDRM definition statement, to specify the TG profile that describes the required characteristics of the VR-based TG. For general information about how an APPN TG is chosen for a given session, see [“Network routing and resource location for APPN nodes” on page 235](#). For details on the operands that you can code in a transmission group profile, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

When a VR-based TG has been selected for a session, the mode name for the session is used to select the subarea Class of Service, which is then used to determine which underlying VRs can be used for the session. Therefore, it is necessary to ensure that VRs and ERs are defined for every subarea Class of Service that might be selected when a VR-based TG is included in the route that is calculated during APPN session establishment.

VR-based TG recommendations

When defining the CP-CP sessions in an APPN network, it is important to remember that APPN networks do not require meshed CP-CP sessions the way subarea networks require meshed SSCP-SSCP sessions. In APPN networks, direct CP-CP sessions between all pairs of network nodes are impractical because of the number of network nodes, and it is not possible in many cases because CP-CP sessions are allowed only with physically adjacent nodes. Instead, only a path of active CP-CP sessions between the origin and destination nodes is needed to deliver a session establishment request, determine an appropriate session route, and establish the session.

In fact, too many CP-CP sessions in an APPN network can cause network nodes to spend too much time processing duplicate requests. For example, when network topology database updates and broadcast search requests are performed, every network node must send or receive every request over every pair of CP-CP sessions that it has with other network nodes. (This ensures that every network node receives the request.) For this reason, establishing CP-CP sessions with every adjacent node might also prove to be undesirable.

At the same time, not enough CP-CP sessions in an APPN network can cause portions of the network to become disjoint, so that CP-CP sessions fail and network topology database updates and broadcast search requests can no longer be delivered to every network node. Therefore, you need to consider (which takes into account network availability requirements) the number of CP-CP sessions and the placement of those sessions. The ideal configuration is one that minimizes the number of CP-CP sessions while still ensuring that, at any time, every network node is logically connected to every other network node by a path of active CP-CP sessions.

The same considerations are necessary when defining VR-based TGs and CP-CP sessions through a subarea network. The following recommendations can simplify the task of determining how many VR-based TGs and CP-CP sessions should be defined, and where in the network they should be located.

Physical adjacency

VR-based TGs and CP-CP sessions should be defined only between physically adjacent VTAM nodes (or composite network nodes). They should not be defined through an intermediate VTAM. For example, in [Figure 30 on page 82](#), VR-based TGs and CP-CP sessions should be defined only between HOST1 and HOST2 and between HOST2 and HOST3. CP-CP sessions between HOST1 and HOST3 (through HOST2) are unnecessary, because session setup flows can be sent from HOST1 to HOST3 using HOST2 as an intermediate APPN node. Also, because RSCV pruning is performed, a two-hop VR-based TG route from HOST1 to HOST2 to HOST3 is pruned to a one-hop VR-based TG route from HOST1 directly to HOST3. This pruning allows a VR-based TG to be used between HOST1 and HOST3, even though it was not explicitly defined.

CMC connectivity

If the subarea network is set up as a communications management configuration (CMC) with multiple VTAMs acting as CMCs, there must be enough VR-based TGs and CP-CP sessions defined between the CMCs to satisfy network availability requirements, based on the considerations described earlier. Although some redundant VR-based TGs and CP-CP sessions might have to be defined, to provide connectivity to all network nodes when TGs carrying CP-CP sessions fail abnormally, meshed VR-based TGs and CP-CP sessions between every pair of VTAM nodes will rarely be required, even when every VTAM is physically adjacent to every other VTAM.

Migration data host connectivity

Migration data hosts are VTAM end nodes that also support CDRM sessions. As end nodes, migration data hosts require the services of a network node server in order to establish sessions with resources in other nodes. The network node server for VTAM end nodes is typically a VTAM network node or composite network node, possibly a CMC host.

If all of the migration data hosts are served by the same network node server, then VR-based TGs and CP-CP sessions need to be defined only between each migration data host and its network node server

(not between pairs of migration data hosts), as shown in Figure 31 on page 84. In some cases, a session is requested between two migration data hosts, and the route is calculated as a two-hop VR-based TG route through the network node server. RSCV pruning reduces this route to a one-hop VR-based TG route directly between the two migration data hosts, and an appropriate VR is selected for the session.




Figure 31. VR-based TGs in a communication management configuration

The same VR selection process is used when the migration data hosts are served by different network node servers and there is a VR-based TG path between the two network node servers. In this case, each migration data host should define VR-based TGs to all potential network node servers. In some cases, a session is requested between two migration data hosts, and the route is calculated as a two-hop VR-based TG route through the network node server. RSCV pruning reduces this route to a one-hop VR-based TG route directly between the two migration data hosts, and an appropriate VR is selected for the session.

If two migration data hosts are served by different network node servers and no VR-based TG path exists between the two network node servers, each migration data host should define VR-based TGs to each adjacent VTAM network node and migration data host to improve the chances of finding a route between the migration data hosts that consists of entirely VR-based TGs.

Selecting the network node server for end nodes

You might want to control which network node is selected by an end node as its network node server (NNS). For example, you might want to exclude a particular network node from network node server responsibilities because acting as a server involves some overhead, such as originating search requests or executing domain broadcasts. You might also want to specify which network nodes should be allowed as backup network node servers in the event the primary network node server becomes unavailable.

You can use a network node server list to indicate the network nodes that can act as servers for an end node. An end node detects adjacent network nodes capable of acting as its network node server through link activation or system definition. If you do not use a network node server list, the adjacent network node chosen as an end node server is directly related to when the end node becomes aware of that particular network node, relative to other adjacent network nodes. Without a network node server list, an end node establishes a CP-CP session with the first such network node that it becomes aware of, and this network node acts as the end node server. There is no predictability in this scheme; it is all timing-dependent. Therefore, if you want the network node server list to be used, you need to activate the list before activating links to network nodes. The configuration list should have the network node server list as one of the first resources to activate.

Initially, a CP-CP-capable physical connection might be temporarily unavailable between an end node and the network node that you prefer to have as the end node primary network node server. This situation requires the end node to establish CP-CP sessions with a backup server specified by the network node server list or, if no network node server list has been defined, with the first network node it can activate a CP-CP capable physical connection with.

By specifying the name of the end node preferred network node server on the end node NNSPREF start option, you allow the end node to automatically switch CP-CP sessions from its current NNS to the preferred NNS when a CP-CP-capable connection to the preferred network node server becomes available.

Creating a network node server list

To restrict the network node server to one that has the same level of support for SLU-initiated sessions as the end node, you can use the SLUINIT operand on the NETSRVR definition statement. For example, to require a network node server to be able to support SLU-initiated sessions, you can define SLUINIT=REQ. The SLUINIT=REQ specifies that CP-CP sessions can be established only with a network node that supports SLU-initiated sessions. SLUINIT=REQUIRED is the default. If you define SLUINIT=OPT, CP-CP sessions are established with a network node server regardless of whether the network node supports SLU-initiated sessions.

You can create a prioritized network node server list by allowing the ORDER operand to default to FIRST on the NETSRVR definition statement. ORDER=FIRST causes the end node to always start at the top of the list when searching for a network node server. If ORDER=NEXT, the end node selects the network node in the list following the last node selected. ORDER=NEXT gives no preference to any node on the list.

To create a network node server list for an end node, create a VTAMLST member containing a VBUILD TYPE=NETSRVR definition statement, and one or more NETSRVR definition statements for each network node that you want in that network node server list. You can also include a nameless entry to allow the end node to select any other known adjacent network node that meets the defined criteria as its network node server. A sample network node server list follows.

```
VBUILD TYPE=NETSRVR
CP1A  NETSRVR  NETID=NETA
CP2A  NETSRVR  NETID=NETA
CP3A  NETSRVR  NETID=NETA
      NETSRVR
```

Note: Any network node named on a NETSRVR definition statement previously considered and rejected as a network node server is not retried during nameless entry processing.

Activating, replacing, and displaying a network node server list

You can activate a network node server list at startup by specifying the appropriate VTAMLST member name in your configuration list, or by using the VARY ACT command at an end node. If you use the VARY ACT command to activate the network node server list after a CP-CP session has already been established, the list is not used until the existing CP-CP session between the end node and its current network node server is lost or terminated. Use the VARY TERM command at the end node to deactivate the CP-CP session between the end node and the network node acting as the end node server. VTAM will then use the network node server list in selecting a new server for the end node.

You can also dynamically replace an existing network node server list. To do this, create a new list and, as described above, use the VARY ACT command to activate it and the VARY TERM command to force VTAM to use the new list.

To display server information for end nodes, use the DISPLAY NETSRVR command at an end node to display the end node network node server list, the end node current network node server, and other network nodes known to the end node.

Using the NNSPREF start option

The NNSPREF start option gives the end node the ability to dynamically switch CP-CP sessions from the current network node server (NNS) to a different (preferred) NNS. This switch occurs automatically in either of the following circumstances:

- A physical connection is successfully established between the end node and the preferred NNS.
- The end node operator manually modifies the NNS preference from the name of one adjacent network nodes to the name of another adjacent network node (NN). In this case, a physical connection to the NN that is taking over the NNS function for the end node must exist.

If you do not specify NNSPREF, or specify NNSPREF with a value of NONE (the default), dynamic NNS switching is not enabled. You can modify NNSPREF using the MODIFY VTAMOPTS command. The value specified on MODIFY NNSPREF effectively gets dynamically added to the existing NETSRVR list table that is active. (To display the value of NNSPREF, use the D NET, VTAMOPTS,OPTION=NNSPREF command.) For more information about the DISPLAY VTAMOPTS command, see [z/OS Communications Server: SNA Operation](#).

Chapter 5. Connecting a subarea node to VTAM

If a VTAM is connected to a subarea node, a physical link between the nodes must be connected and functional at both ends. A subarea node can be a host processor (containing another VTAM) or a communication controller (containing an NCP).

The following list shows where to find information about the following connections:

- [“Connecting two VTAMs using channels” on page 87](#)

Note: Earlier releases of VTAM support type 2.1 casual connection, a way of connecting two VTAMs in a peer-to-peer relationship using low-entry networking (LEN). While type 2.1 casual connection is still supported, IBM recommends that you use APPN to connect two VTAMs in a peer-to-peer relationship. See [Chapter 4, “Connecting an APPN node to VTAM,” on page 39](#) for information about connecting two VTAMs using APPN.

Connecting two VTAMs using channels

The following channel connections can be used to connect two VTAM subareas:

- [“Channel-to-channel adapter connection” on page 87](#)
- [“Multipath channel connections” on page 89](#)

Channel-to-channel adapter connection

A channel-to-channel adapter connection between two host processors is defined using a channel-attachment major node. You define two channel-attachment major nodes for each connection. Each channel-attachment major node represents each VTAM's view of the connection. A GROUP definition statement defines the type of links that follow and also carries other operands that sift down to following definition statements that do not explicitly override them.

Code one LINE definition statement for each channel-to-channel adapter. The LINE definition statement defines to VTAM the characteristics of the line side of the adapter. Code one PU definition statement for each LINE definition statement.

The definition statements in the following example define the channel link between the two VTAMs shown in [Figure 32 on page 88](#). On the GROUP definition statement, the missing interrupt handler (MIH) operand is coded with a value of YES so that the channel link becomes inoperative after the time period specified on the REPLYTO operand expires. (Otherwise, the channel link appears operative, but VTAM cannot use it.) The default for the MIH operand is NO.

The LINE definition statement is used to specify the channel unit address (ADDRESS=51A) used in VTAM at one end of the channel link.

You can use the following operands on the LINE definition:

- MAXBFRU defines the number of 4 KB pages of storage that are used to buffer PIUs for transmission over the channel link. The PIUs are blocked into the 4 KB pages, transferred over the channel, and then unblocked into VTAM I/O buffers in the other VTAM.
- DELAY can be used to delay the operation of the data transfer so that more PIUs can be buffered and transferred in a single channel I/O operation.

For more information about selecting values for these operands, see [Chapter 20, “Tuning VTAM for your environment,” on page 499](#).

Though the PU definition statement identifies the physical unit type at the other end of the connection to be an NCP (PUTYPE=4), the channel link actually connects to another VTAM, a physical unit type 5. The

TGN operand specifies a transmission group number for the channel-link transmission group. Because parallel transmission groups between two subareas require unique TG numbers, the TGN operand is used.

TSTCAX	VBUILD	TYPE=CA	CHANNEL-TO-CHANNEL (CTC) CONNECTION
*			
CTCGRPX	GROUP	LNCTL=CTCA, MIH=YES, REPLYTO=10.0, : ISTATUS=ACTIVE	CTC CONNECTION LINK INOPERATIVE - START I/O TIMEOUT CHANNEL PROGRAM TIMEOUT VTAM INITIAL STATUS
*			
CTCSA51L	LINE	ADDRESS=51A, DELAY=.1, MAXBFRU=8, : ISTATUS=ACTIVE	CHANNEL UNIT ADDRESS CHANNEL DELAY TIMER 4K PAGE BUFFER FOR CTC DATA TRANSFER VTAM INITIAL STATUS
*			
CTCSA51P	PU	PUTYPE=4, TGN=3, : ISTATUS=ACTIVE	PHYSICAL UNIT TYPE = VTAM TRANSMISSION GROUP NUMBER VTAM INITIAL STATUS

The capability of having parallel transmission groups allows you to add multiple channel attachments to an adjacent subarea host. With parallel transmission group support, VTAM permits the transmission group number to be specified on the PU definition statement of the channel-attachment major node and on the PATH definition statement. For information about how this affects network routing, see [“Parallel sessions using parallel transmission groups”](#) on page 273.

Although you can have as many as 255 transmission groups, only 16 of these can be defined because the maximum number of explicit routes that can be defined between two adjacent VTAMs is 16. Each transmission group is single-link capable only. [Figure 32 on page 88](#) illustrates a possible configuration of parallel transmission groups.

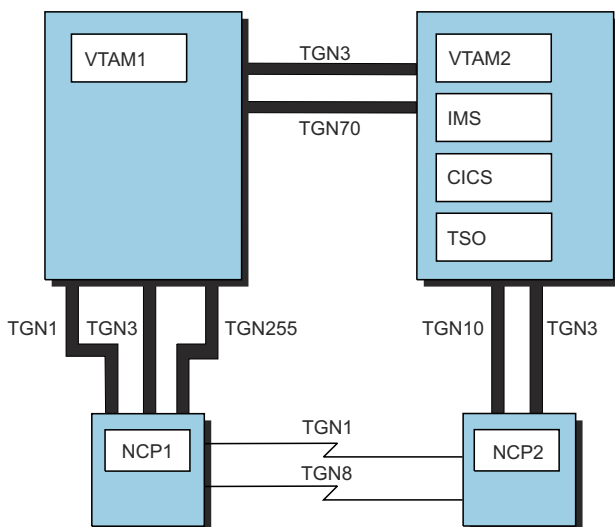


Figure 32. Parallel transmission groups in multiple domain environment with NCP

Two channel-to-channel connections are configured in [Figure 32 on page 88](#). A LINE and a PU definition statement are required for each of these channel connections between the two VTAMs. The TGN operand for one of the channel connections specifies a transmission group number of 3, as shown in the previous example. The LINE definition statement that defines the other channel-link connection must have a TGN operand that specifies a transmission group number of 70. Parallel transmission group support limits each transmission group to a single channel-to-channel link. The two channel-link connections (TG3 and TG70) between the VTAMs cannot be defined as a single logical transmission group composed of the two channel links.

Multipath channel connections

Multipath channel (MPC) connections allow you to code a single transmission group (TG) for host-to-host communication that uses multiple write-direction and read-direction subchannels. The subchannels can be assigned to one or more physical channels based on how you have defined your I/O subsystem to the operating system. Because each subchannel operates in only one direction, the half-duplex turnaround time that occurs with other channel-to-channel connections is reduced.

To define an MPC connection between two subarea nodes, code a channel-attachment major node (VBUILD=CA) with LNCTL=MPC on the GROUP definition statement. For more detailed information about defining subarea MPC connections, see [z/OS Communications Server: SNA Resource Definition Samples](#).

If you code a TG in which the subchannels are divided between two physical channels, you can increase availability because the TG will have a path to use, even if one physical channel is down. Because each TG can use more than one channel, and because the turnaround time required for half-duplex is reduced, throughput is increased.

The subchannels on the physical channel are represented by the subchannel addresses coded on the READ and WRITE operands on the LINE definition statement in the channel-attachment major node. One READ subchannel in one host and the corresponding WRITE subchannel in the other host form a complete path. The READ subchannel address and the corresponding WRITE subchannel address must reference the same physical connection between the two nodes; the two addresses do not need to be identical. You should consider the ratio of inbound and outbound traffic for a host in deciding how many write subchannels and read subchannels you want to code.

If multiple write subchannels are used, traffic is distributed among the subchannels as traffic requests increase. Arriving data is resequenced before being presented to the application. Priority is assigned to the write subchannels in the order that they are defined.

If you want two TGs, with each TG dividing its subchannels across two parallel physical channels, you have to code two groups, one for each TG. The LINE definition statement for each TG will include subchannel addresses on the READ and WRITE operands from each of the two physical channels.

If a subchannel becomes inoperative, the TG remains active as long as one subchannel is available in each direction. If the last subchannel in either direction becomes inoperative, issue a VARY INACT command, correct the problem (all subchannels must be allocatable when activated or activation will fail), and then issue a VARY ACT command to reactivate the major node.

Note: MPC connections in an APPN network enable two nodes to communicate using APPN protocols over MPC connections. For a description of MPC connection support in an APPN network, see [“Multipath channel connections”](#) on page 40.

Chapter 6. Using Enterprise Extender (EE)

Companies continue to rely on older applications residing on mainframes. Initially, access to these applications was through SNA. Today, newer applications are generally based on TCP/IP. Initially, companies supported both these networks separately, but are now seeking ways to consolidate their SNA traffic onto the TCP/IP network. However, it is not economically practical to convert existing SNA applications to TCP/IP and, in many cases, conversion might even be technically impractical because of the lack of source code, adequate skills, or both.

With Enterprise Extender (EE), you can extend the reach of SNA applications and data to include TCP/IP networks and IP-attached clients with levels of reliability, scalability, and control similar to those that SNA users have used. EE integration uses standard IP technology and does not require new hardware or new software in the IP backbone.

You can use an IP network for SNA sessions with Enterprise Extender, which provides enablement of IP applications and convergence on a single network transport while preserving SNA application and endpoint investment. An EE connection is a logical connection that represents IP connectivity from a host to a specified IP address or host name. Conceptually, an IP network looks like an APPN/HPR transmission group (TG) in a session route.

In summary, EE is an extension of High Performance Routing (HPR) technology that provides encapsulation of SNA application traffic within UDP frames by HPR-capable devices at the edges of an IP network. To the IP network, the SNA traffic is UDP datagrams that get routed without hardware or software changes to the IP backbone. To the user, the session is normal SNA with predictable performance and high availability. By wrapping the SNA application in this way, EE enables SNA data to be carried over an IP backbone without changing either the SNA applications or the IP hardware.

This information is organized into the following sections:

- [“Overview” on page 92](#)
- [“Designing the EE network” on page 94](#)
- [“Configuring the EE network” on page 99](#)
- [“Configuring the EE connection network” on page 116](#)
- [“EE security considerations” on page 127](#)
- [“Tuning the EE network” on page 130](#)
- [“Advanced coding considerations for EE” on page 137](#)
- [“Troubleshooting EE problems” on page 152](#)

In addition, this topic contains information about various topics including:

- [“EE connection network reachability awareness ” on page 137](#)
- [“Benefits of defining multiple Enterprise Extender virtual routing nodes” on page 122](#)
- [“Parallel transmission groups \(TGs\) ” on page 96](#)
- [“TCP/IP MTU size for EE ” on page 144](#)
- [“Running EE in constrained or virtualized environments” on page 145](#)
- [“RTP transmission stall operator awareness and recovery support” on page 146](#)
- [“Load balancing” on page 147](#)
- [“Transmission group profiles” on page 147](#)

For more information about Enterprise Extender, see *Migrating Subarea Networks to an IP Infrastructure*, SG24-5957-00 (a member of the IBM Redbooks library) <http://www.redbooks.ibm.com>.

Overview

This section provides information about the following topics:

- [“Benefits of Enterprise Extender” on page 92](#)
- [“Availability of Enterprise Extender” on page 92](#)
- [“Hardware requirements” on page 92](#)
- [“EE reliability and strategy” on page 93](#)
- [“Using EE and extended border node \(EBN\) as a replacement for SNI” on page 94](#)
- [“EE implementation considerations” on page 94](#)

Benefits of Enterprise Extender

EE provides the following benefits:

- SNA transport over a native IP network with no changes required to SNA applications
- SNA application connectivity using an IP backbone that supports preservation of SNA transmission priority
- When used with an extended border node, EE provides a mechanism to replace SNA Network Interconnection (SNI) function in a way that does not require SNA legacy hardware, such as the 3745. See [“Using EE and extended border node \(EBN\) as a replacement for SNI” on page 94](#) for more information.
- SNA traffic can use OSA Ethernet (EE can use any z/OS-supported IP network connection)
- Support for IPv4 and IPv6 addressing models
- End-to-end failure protection and data prioritization
- Compatibility with IPsec and SNA session-level encryption

Requirement: EE is valid only for APPN configurations and is an extension of APPN and high performance routing (HPR) protocols. Your subarea network must be migrated to APPN for the subareas where EE is to be deployed. For more information about APPN, see [Chapter 2, “VTAM networking concepts,” on page 3](#) and [Chapter 4, “Connecting an APPN node to VTAM,” on page 39](#). Also see *Migrating Subarea Networks to an IP Infrastructure* SG24-5957-00.

Availability of Enterprise Extender

Enterprise Extender can be implemented on the following platforms:

- Communications Server NT
- Communications Server for Linux
- Communications Server for AIX®
- Communications Server for OS/2
- z/OS Communications Server
- Personal Communications
- Microsoft Host Integration Server
- Cisco SNASw

Restriction: EE is not supported by VTAM on any VM or VSE platform.

Hardware requirements

EE can use any interface supported by the TCP/IP stack. Use OSA and HiperSockets for optimal performance and function. With OSA, you can use inbound workload queueing to further improve throughput for both inbound Enterprise Extender (EE) packets and for inbound non-EE traffic over the same interface. For more information, see [QDIO inbound workload queueing in z/OS Communications](#)

Server: IP Configuration Guide. For intra-CEC communication, EE-over-HiperSockets provides superior performance unless CPU availability is limited. If CPU availability is limited, then EE communication using a shared OSA adapter provides optimal performance.

EE reliability and strategy

Because EE relies on IP network strategy and integrity, effective IP network design is essential to ensure successful EE implementation. The robustness of EE depends on the stability of your IP network configuration, just as any mission-critical TCP application does. With a stable IP network configuration in place, EE ensures the availability of SNA applications.

In planning your EE environment, discuss your configuration with the WAN environment architect to determine whether VPN or your firewall will be affected. You should also communicate with those responsible for MVS, because EE can affect the way OSA, OSAE, or CIP are defined.

Base a well-designed IP Infrastructure on the following considerations:

- The dynamic routing update protocols being used (such as OSPF, RIP, RIPv2, EIGRP, and so on)
- Use of a private intranet in contrast to the public Internet
- The type of network interfaces used, and MTU size
- The IP security mechanisms (NAT, IPSec, firewall)
- The quality of service (QoS) for the WAN

Figure 33 on page 93 shows the differences between a SNA network and EE using an IP network.

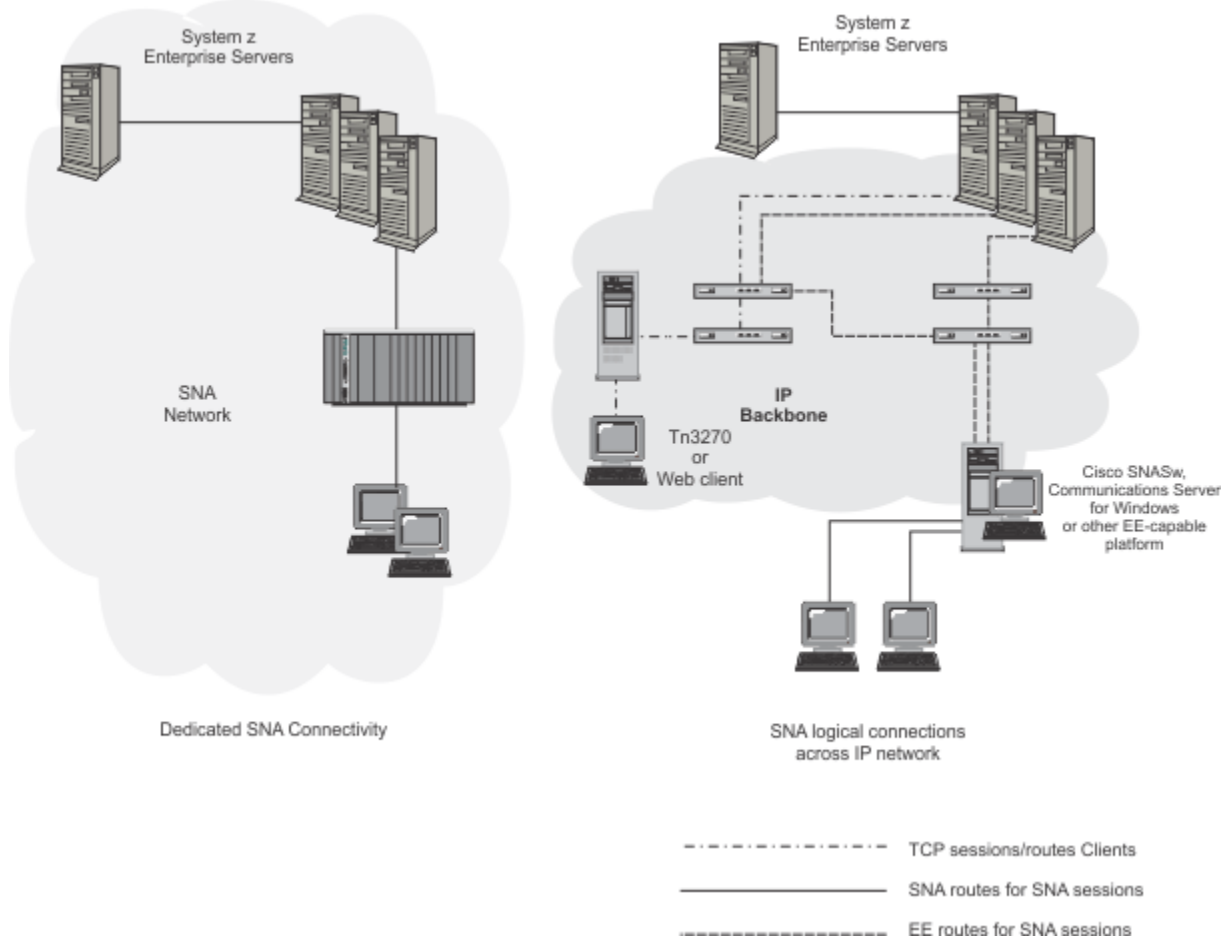


Figure 33. Comparison between an EE network and a SNA network

Using EE and extended border node (EBN) as a replacement for SNI

Enterprise Extender used with the APPN extended border node function can replace SNA Network Interconnection (SNI) function in a way that does not require SNA legacy hardware, such as the 3745. [Figure 34 on page 94](#) shows how EE and EBN work together. Because EE with EBN requires that both endpoints be migrated from SNI, using EE with EBN as an SNI replacement technology requires cooperative changes with other partner companies.




Figure 34. How EE and EBN work together

EE implementation considerations

Before implementing EE, carefully consider your business requirements, especially the following questions:

- Identify your preferences; will you need dynamic or predefined connections?
- Evaluate your priorities. Which of the following considerations is most important in your environment?
 - Less maintenance
 - Less definition
 - More control
 - More security
- Do you want to use host names or IP addresses for IP configuration? Do you maintain a DNS to resolve host names to IP addresses? The benefit of maintaining your own DNS is that you can keep all IP addresses in a central location. If the IP addressing scheme changes, you do not have to update VTAM definitions because DNS resolves the host names to the correct IP addresses.
- Decide whether you want to use EE connection network and which EE partner nodes will participate in the EE connection network.
- Determine whether network address translation (NAT) is being used in your IP network. See [“Network address translation \(NAT\) considerations” on page 101](#) for more information related to using host names.

Designing the EE network

Proper design is critical as you plan how to implement EE in your network. To make your implementation effective, consider these design guidelines:

- Spread EE connections for a balanced workload and avoid sending all your EE connections through your primary host.
- Use direct EE connections to endpoints.
- Consider an EE connection network, which might simplify definitions.
- Avoid using VTAM as a router.

This section explains how static VIPA addresses represent EE endpoints to z/OS Communications Server, how you can use multiple VIPA addresses to your advantage, how you need to consider class of service (COS) and quality of service (QoS), and how to decide whether you should specify addresses using host name definitions or IP address definitions.

Distinctions between an EE network and an EE connection network

You can implement EE in one of two ways:

- A simple EE connection. In this implementation there are individual connections between endpoints. This connection can be either static with each endpoint of the connection defined with a predefined switch PU or dynamic with the DYNPU operand allowing for dynamically defined dial-in connections.
- An EE connection network. This implementation uses a shared access transport facility (SATF) involving a common virtual routing node for communication. This connection is dynamically built as needed. EE connection network is explained in more detail in the section [“Configuring the EE connection network”](#) on page 116.

Evaluate whether an EE connection network applies to your network. A connection network is a representation of a shared access transport facility (SATF) that enables nodes to identify their connectivity to the SATF by a common virtual routing node in order to communicate without having individually defined connections to one another.

Many customers prefer to use an EE connection network because it provides a balance between control and usability. You do not need an EE connection network to have dynamic definitions; but if you use dynamic definitions, an EE connection network can be a simplified means of enabling dynamic connectivity between a set of EE partners.

Generally, all PUs associated with an EE connection network environment are dynamic. There are no predefined EE partner addresses, just the VRN definitions. Each node must have a predefined APPN connection (EE or otherwise) that provides end-to-end APPN locate capability in addition to the VRN definitions. With an EE connection network, the following are reduced:

- Definitions
- Activation actions
- Maintenance of switched major nodes

Characteristics of EE connections

EE is a simple set of extensions to the existing, open HPR technology. It integrates the HPR frames using User Datagram Protocol/Internet Protocol (UDP/IP) packets to deliver dependable SNA networking benefits to corporate intranets and the Internet. EE provides the following benefits:

- Ability to use all z/OS Communications Server IP DLCs
- Seamless IP/SNA integration

For more information about HPR, see [Chapter 15, “Implementing an APPN network,”](#) on page 385.

From the perspective of the HPR network, the connection across the IP backbone looks like a logical link; in the IP network, the SNA traffic is just UDP application datagrams that are routed to the IP backbone. Unlike gateways, there is no protocol transformation; and unlike common tunneling mechanisms, the integration is performed at the routing layers without the overhead of additional transport functions.

Sessions do not have to go through a specific subarea (VTAM and NCP) attachment point (boundary function) because EE is based on peer networking. This helps to optimize network flow.

Restriction: Because EE relies on the APPN/HPR protocols, only SNA sessions are supported. SNA sessions involving both independent LUs and dependent LUs are supported across EE connections, with DLUR providing the boundary services for the dependent LUs.

Static VIPA considerations

EE endpoints are represented to the z/OS Communications Server network using static VIPA addresses. An EE endpoint can have one or more EE static VIPA addresses. A VIPA address can be dedicated to an EE endpoint, or it can be shared with other applications.

Guideline: Use a dedicated EE VIPA address that is not shared with other applications to simplify problem determination and network management.

Restriction: EE does not support dynamic VIPA.

Using multiple VIPAs

There are several advantages to having multiple VIPAs.

- You can provide different VIPAs to different vendors, which eases the tasks of network management, problem determination, and firewall administration.
- You can define multiple connection networks.
- There is more potential for flexibility in making IP routing decisions.

Requirement: All local Enterprise Extender VIPAs must belong to the same TCP/IP stack. VTAM establishes affinity to a single TCP/IP stack when the first EE line is activated; VTAM cannot use multiple stacks concurrently.

Parallel transmission groups (TGs)

You have the option of defining parallel EE TGs by using different EE VIPAs for one or both of the endpoints. This method of defining parallel TGs might be useful in some cases, because, depending on the IP network topology and the IP routing protocol in use, different EE TGs can map to different IP routes.

Rule: Do not code different SAP values to define parallel TGs.

Rule for multiple EE connections

z/OS Communications Server restricts how multiple EE connections are established between two EE endpoints.

When establishing multiple EE connections between the same two EE endpoints (parallel EE connections), you cannot use the same IP address pair for more than one EE connection. Instead, every EE connection must use a unique IP address pair. You can define parallel EE connections by using a different local static VIPA address or remote IP address (or both) for each connection. z/OS Communications Server does not prevent parallel EE connections from being established as long as each connection uses a unique IP address pair. There is no limit to the number of EE connections that can be established using unique IP address pairs.

Before V1R8, an alternate method allowed parallel EE connections between the same IP address pair. As of V1R8, the initiation of parallel EE connections between the same IP address pair is no longer supported. This restriction applies only to parallel EE connections that are initiated by VTAM. An inbound parallel EE connection request, that uses the same IP address pair and different SAP values, will be accepted by VTAM.

Rule: Do not define parallel EE connections between the same IP addresses. Because each of these parallel connections traverse the same IP network, there are no advantages with regards to either bandwidth exploitation or availability.

IP multipath considerations

The IPCONFIG MULTIPATH and IPCONFIG6 MULTIPATH parameters in the TCP/IP profile enable the multipath routing selection algorithm for outbound IP traffic (the default value is NOMULTIPATH.). When multipath routing is enabled, there are two options: PERCONNECTION and PERPACKET; PERCONNECTION is the default value. If multipath routing is enabled in the TCP/IP stack for EE traffic, it does not matter which option you choose. In either case, IP uses a per-batch of packets from VTAM approach. Because there is no UDP connection for EE, true per-connection multipath is not possible. Per-packet multipath routing is too granular because it leads to too much resequencing overhead at the RTP receiving endpoint.

Multipath for Enterprise Extender (EE) is disabled by default. Enabling multipath for EE can lead to poor performance. A common problem is that one or more of the routes might not have a path to the destination IP address. This can result in packet loss of up to as much as 50 percent. Another problem is that one route might be slower than the others, which can cause half of the packets to arrive out of order. Both of these conditions result in poor send rates for the HPR connections that are using the EE connections. Multipath for EE is controlled with the VTAM start option MULTIPATH. To enable multipath for EE, code MULTIPATH in the TCP/IP profile and code MULTIPATH=TCPVALUE in the VTAM start options. To disable multipath for EE, code the VTAM start option MULTIPATH to the value NO or allow MULTIPATH to default. If MULTIPATH was coded in the IPCONFIG statement, all other applications will continue to use multipath.

Class of Service preservation dependencies

With EE, you can adhere to SNA transmission priorities across IP networks, sending batch transfers and smaller interactive sessions efficiently. EE maps SNA transmission priorities to the IP type of service (ToS) header. To take advantage of the mapping, you must correctly configure IP routers and switches to accommodate the IP ToS settings. If the configuration is incorrect, SNA priorities are not applied to data as it travels across the IP network.

Distinct routing algorithms can differentiate among IP packet priorities, and these algorithms have to be established as you set up routing hardware. You need to set up routers to process the ToS value that you want. SNA transmission priorities are competing against native IP applications and every other application using the WAN; for example, FTP, Telnet, and web-based traffic. EE is just another application competing for priority.

Establish traffic prioritization (ToS settings) as you design your EE network. Consult the appropriate documentation specific to your router to establish the correct traffic prioritization. Host personnel and network engineers need to agree on ToS values so that applications get the priority they need. If this is not done before implementation, EE will not be effective and SNA traffic will not get the proper priority on the router. [Figure 35 on page 98](#) shows how ToS settings affect IP traffic.




Figure 35. How ToS settings affect IP traffic

See [z/OS Communications Server: IP Configuration Guide](#) for more information about ToS.

Comparing host name and IP address definitions

To configure EE, the IP addresses on both sides of a connection need to be known; one on the local side, and one on the remote side. EE creates the logical connection between these two addresses. The following rules apply to using a host name instead of an IP address.

Rules:

- If you use EE across IPv6 networks, you must use host names to define endpoints.
- If you define EE connection networks, use host name definitions instead of IP addresses.
 - If you are using an EE connection network over an IP network with NAT boundaries, you must use host names.
 - If you are using an EE connection network over an IP network without NAT boundaries, using host names ensures compatibility with future IP addressing schemes.

Types of definitions

When defining EE connections, you can choose between predefined (static) or dynamic definitions.

Static definitions are more secure, which is often a priority for host connections. The disadvantage is that more labor is required to maintain definitions.

Dynamic definitions are less secure but require fewer definitions to maintain.

<i>Table 5. Comparison between static and dynamic definitions</i>	
Predefined (static) definitions	Dynamic definitions
Elaborate implementation; many individual definitions; DYNPU=NO	Streamlined implementation; no individual definitions; DYNPU=YES
More secure	Less secure
Ability to manage specific connections with user-defined names	Minimal control over assigned names
More flexible - different characteristics for different connections	Less flexible - most connections have same set of characteristics
Transmission group number (TGN) helps identify partner; specific TGNs can differentiate media type to enhance problem determination	Control of TG number only by defining TGN on the DYNTPU=EE PU entry in a model major node
More resources (definitions) to maintain	Self-managing

It is possible to have EE connections between the following endpoints:

- Two static end points
- One static end point and one dynamic end point using DYNPU=yes
- Two dynamic end points defined with connection network

Configuring the EE network

You need to plan how you want to configure the EE network before you start configuring.

Procedure

Take the following steps to plan how to configure the EE network:

1. Draw a logical picture of your network and decide what types of EE connectivity you want to implement. Four choices are represented in [Figure 36 on page 100](#).

- Host-to-host with same NETID

This type of EE connection is between z/OS hosts within your data center.

- Host-to-host with extended border node to vendor

This type of EE connection uses EE along with VTAM extended border node to connect two vendors, replacing existing SNI connectivity.

- Host-to-branch

This type of EE connection sets a path to one or more branches across a WAN. This connection requires an EE-capable access point in the branch such as IBM Communications Server for Linux on Intel or PCOMM.

- HiperSockets within a CEC

This type of EE connection is made between LPARs within a CEC.




Figure 36. Four types of EE connectivity

2. Decide on the level of control, flexibility, and ease of use you need, and a naming scheme.
See [“Types of definitions”](#) on page 98.
3. Review the following terms:

model major node

Defines a model PU to be used for dynamically created EE PUs on inbound connections.

switched major node

Defines the switched TGs with the following statements:

- The PU statement defines the remote node.
- The PATH statement defines the remote IPv4 address, or a host name that can be resolved into an IPv4 or an IPv6 address.

XCA major node

Defines the IP port with the following statements:

- The PORT statement identifies the name of a port through which an HPR connection through the IP network is made. The PORT statement must be set to MEDIUM=HPRIP for EE.
- One or more GROUP statements define a set of lines and PUs that represent potential EE partners. All lines and PUs under a given group are associated with a particular local IP address or host name.

4. Decide what types of EE connectivity you want to implement.

All four types have these benefits:

- Increased bandwidth because EE can use OSA adapters.
- Elimination of dependency on SNA legacy hardware such as the 3745 and ESCON-based channel connectivity.

In addition, for host-to-host with EBN, another benefit is the elimination of Gateway NCP definitions. With HiperSockets, EE provides SNA applications with the ability to communicate across a high-speed low-latency data communication path, with no network connection and no special hardware.

Results

Requirement: To fully evaluate the best option for your environment, you must ensure that your WAN hardware and firewalls accommodate the EE connectivity of your choice. Any hardware device performing IP filtering must allow UDP traffic on ports 12000 through 12004 both inbound and outbound. If your organization does not want to allow UDP packets to flow through the firewall, you can limit UDP traffic to a subset of trusted partners using IP filtering.

Rule: Do not modify the default value for the IPPORT operand or EE interoperability problems might result.

Remember that EE works no better than your IP network. A healthy WAN environment and stable IP connectivity ensures successful EE connections.

Preservation of SNA transmission priority

In Logmode/COS assignment, every session is assigned an APPN COS with an associated transmission priority so that data flowing over that session is one of the four transmission priorities (network, high, medium, or low). EE attempts to preserve that transmission priority across the WAN by mapping the transmission priority to an appropriate IP type of service (ToS) byte setting and UDP port. Configure IP WAN routers to handle appropriate transmission priorities.

Guideline: EE also attempts to preserve SNA transmission priority by mapping each SNA TP to a corresponding UDP port. However, traffic prioritization using ToS byte management is generally preferable to port-based prioritization, because port-based prioritization schemes do not work well in the presence of IP fragmentation or IPSec-encrypted traffic.

Network address translation (NAT) considerations

Because EE depends on IP to establish its connectivity, NAT can pose difficulties in some situations. If you choose to use NAT, EE requires static NAT. Dynamic NAT is not appropriate because an address needs to be determined at configuration time. If a NAT boundary is associated with an EE connection network path, then host name-based EE definitions are required. See [“Configuring the EE connection network” on page 116](#) for more information.

Restrictions:

- Do not use Network Address Port Translation (NAPT) if you are using EE.
- Do not use dynamic NAT if you are using EE.

Steps for configuring and activating an EE network

The following procedure is an overview of the steps required for configuring and activating an EE network:

To configure an EE network, take the following steps:

1. [“Step 1: Prepare your TCP/IP configuration” on page 102](#)
2. [“Step 2: Evaluate IP address resolution” on page 103](#)
3. [“Step 3: Prepare your VTAM definitions” on page 105](#)

To activate an EE network, see [“Activating EE” on page 109](#).

Step 1: Prepare your TCP/IP configuration

Procedure

1. Reserve the UDP ports for EE using the PORTRANGE statement or PORT statements to prevent another application from using one of the EE ports.

Tip:

- If you do not reserve the UDP ports, you might have a conflict when another application uses the port. For example:

```
PORTRANGE
12000 5 UDP NET ; RESERVE UDP PORTS 12000-12004 FOR EE
```

- Use the MVS job name associated with the VTAM started task to reserve UDP ports that are to be used for EE network connections. The MVS job name for a given started task can be assigned based on various inputs. These inputs are examined in the following order:
 - a. The job name specified in the JOBNAME= parameter or the identifier specified on the MVS START command. In the PORTRANGE example, VTAM was started with an identifier of NET and therefore the job name specified was NET.
 - b. The job name specified on the JOB JCL statement within the member.
 - c. The member name.

Note: You cannot use wildcard characters when you specify the VTAM job name for EE UDP port reservations on either the PORT statements or PORTRANGE statement.

2. Evaluate your IP routing environment. See [Static VIPA considerations](#) for more information. Also identify the number of static VIPAs.
3. Define one or more static VIPAs for the EE connection.

The following are the TCP/IP profile requirements for Enterprise Extender:

- For IPv4, define the static VIPA addresses for EE by using the DEVICE, LINK, and HOME statements. The following example defines two local static IPv4 VIPA addresses for EE:

```
;*****
; Enterprise Extender VIPAs, IPv4
;*****
DEVICE VIPA01 VIRTUAL 0
LINK LVIPA1 VIRTUAL 0 VIPA01
DEVICE VIPA02 VIRTUAL 0
LINK LVIPA2 VIRTUAL 0 VIPA02
HOME
  92.1.1.1 LVIPA1
  92.1.1.2 LVIPA2
```

- For IPv6, define the static VIPA addresses for EE by using the INTERFACE statement. The following example defines one static IPv6 VIPA address for EE:

```
;*****
; Enterprise Extender VIPA, IPv6
;*****
INTERFACE VIPA6 DEFINE VIRTUAL6 IPADDR 2001:0db8::91:1:1:2
```

- Optionally, enable SOURCEVIPAs. The EE static VIPA is used as the source IP address for all EE traffic, regardless of the SOURCEVIPAs setting. EE uses the static VIPA as the source IP address in all outbound traffic regardless of the SOURCEVIPAs setting. If you have NOSOURCEVIPAs coded or defaulted, non-EE traffic uses the outbound interface address as the source IP address but EE traffic still uses the static VIPA as the source IP address.
4. Define IUTSAMEH connectivity either by using DYNAMICXCF or by creating manual IUTSAMEH definitions in the TCP/IP profile.

If you configure manual IUTSAMEH definitions and also configure the DYNAMICXCF definition, the manual IUTSAMEH definitions will coexist with the other connectivity (XCF and HiperSockets) definitions that dynamic XCF can generate.

To dynamically configure IUTSAMEH, do the following step:

- IUTSAMEH definitions are dynamically defined when dynamic XCF support is enabled using IPCONFIG DYNAMICXCF or IPCONFIG6 DYNAMICXCF. For IPv4, the device name is IUTSAMEH and the link name is EZASAMEMVS. For IPv6, the interface name is EZ6SAMEMVS.

To manually configure IUTSAMEH definitions, do the following step:

- For IPv4, define and start IUTSAMEH using the DEVICE, LINK, HOME, and START statements.
- For IPv6, define and start IUTSAMEH using the INTERFACE and START statements.

Following is an example of a manual configuration:

```
;*****  
; VTAM to TCPIP Stack, IPv4  
;*****  
DEVICE IUTSAMEH MPCPTP  
LINK EELINK MPCPTP IUTSAMEH  
HOME 93.1.1.1 EELINK  
START IUTSAMEH  
  
; *****  
; VTAM to TCPIP Stack, IPv6  
; *****  
INTERFACE EELINK6 DEFINE MPCPTP6  
  TRLENAM IUTSAMEH  
  IPADDR 2001:0db8::91:1:1:1  
START EELINK6
```

Note: If you choose to manually configure the IUTSAMEH connections and the IP address is used only for EE, you do not have to specify the IP address for IUTSAMEH. Warning messages can be ignored.

Tip: To simplify EE definitions, consider defining IUTSAMEH connections using the IPCONFIG DYNAMICXCF statement for IPv4 and the IPCONFIG6 DYNAMICXCF statement for IPv6.

Step 2: Evaluate IP address resolution

Consider the following questions:

- Do you have a DNS infrastructure?

Define name-to-address resolutions in a local hosts file (for example, ETC.IPNODES) or in the appropriate DNS zone files. The following are examples of how to code the resolutions in ETC.IPNODES:

```
92.3.1.1          SSCP2A.RALEIGH.COM  
2001:0db8::91:1:1:1 SSCP2AV6.RALEIGH.COM
```

- Do you use host name support or explicit address coding?

There are several options for specifying the source VIPA address to be used for a specific EE connection:

- Use the TCPNAME start option to identify the TCP/IP stack name from which the IPv4 source VIPA address is to be obtained.
- Use the IPADDR start option to explicitly identify the IPv4 or IPv6 source VIPA address.
- Use the HOSTNAME start option to identify the host name value to be resolved using name-to-address translation into the correct IPv4 or IPv6 source VIPA address.
- Use the HOSTNAME or IPADDR parameters on individual GROUP definition statements within the EE XCA major node.

Guideline: If using the HOSTNAME parameter in a common INET environment, see [TCPNAME](#) in [z/OS Communications Server: SNA Resource Definition Reference](#).

Tip: The HOSTNAME and IPADDR parameters are mutually exclusive. The HOSTNAME parameter value overrides the IPADDR value. See [z/OS Communications Server: SNA Resource Definition Reference](#) for details.

For predefined EE connections, you can specify or obtain the remote IP address of an EE connection in one of the following ways:

- Specify the IPv4 or IPv6 address of the remote partner on the IPADDR operand on the PATH statement of the switched major node used for EE.
- Specify the TCP/IP host name of the remote address. This is specified using the HOSTNAME operand of the PATH statement of the switched major node used for EE. The host name can be resolved to an IPv4 or an IPv6 address.

If you use host name instead of an IP address, you must perform the following additional steps to enable this function:

1. Configure the resolver.

VTAM uses the TCP/IP system resolver to perform name-to-address resolution. The system resolver, in general, first attempts to resolve the name by way of one or more name servers (as defined by resolver configuration statements). If unsuccessful, the resolver then attempts to resolve the name using a local host table, such as HOSTS.SITEINFO (only for IPv4 addresses) or ETC.IPNODES. For more information about defining name servers to the resolver, see [z/OS Communications Server: IP Configuration Guide](#).

The search order for selecting the TCPIP.DATA file to use is documented in [z/OS Communications Server: IP Configuration Guide](#). VTAM uses the native MVS search order. The concept of data set concatenation does not apply here. You may choose to use SYSTCPD for allocating the TCPIP.DATA file. In this case, this statement must be included in your VTAM start procedure.

If you want to use a local table to resolve host names instead of a name server, see [z/OS Communications Server: IP Configuration Guide](#).

Note: Enterprise Extender processing uses the native MVS search order for choosing the correct local host table. See [z/OS Communications Server: IP Configuration Guide](#) for details on the search order.

2. Configure the TCPIP.DATA file.

3. See [z/OS Communications Server: IP Configuration Guide](#) for information about UNIX Systems Services security considerations and required security definitions.

Guidelines:

- While architecturally able to extend to a maximum of 255 characters, the length of the EE-related HOSTNAME operand is restricted by z/OS Communications Server to 64 characters. However, the length of the host name you choose to represent the static VIPA should be much shorter than the maximum. The host name is exchanged between APPN nodes in APPN topology and during APPN session setup flows. Those flows carry the host name within APPN control vectors, which are limited to 255 bytes. Calculations using the contents of the APPN Route Selection Control Vector (RSCV) indicate that limiting the fully qualified host name to 40 characters should allow the host name to be exchanged comfortably within reasonably sized APPN networks.
- If you choose to use the host name function, it is crucial for consistent and predictable EE connection establishment that a given host name resolves to a single IP address. At the TCP/IP stack that owns the host name, the name-to-address resolution should produce the static VIPA address associated with the host name. At a remote EE endpoint wanting to connect to the TCP/IP stack that owns the host name, the name-to-address resolution should produce one of two addresses:
 - If no network address translation is required (for instance, if both nodes are behind the same firewall), the name-to-address resolution at the remote EE partner should generate the static VIPA address of the TCP/IP stack that owns the host name.
 - If network address translation (NAT) is required (for instance, the nodes are separated by a firewall), the name-to-address resolution at the remote EE partner should generate a NAT address.

The NAT address is the address that should point to the static VIPA address at the target TCP/IP stack.

- If you use a name server to resolve host names and your name server runs on a z/OS Communications Server node that uses SOURCEVIPA, code the VIPA address for NSINTERADDR. For more information about SOURCEVIPA, see [z/OS Communications Server: IP Configuration Guide](#).
- If you use a name server to resolve host names and your name server runs on a z/OS Communications Server node that uses VIPA addresses, but does not use SOURCEVIPA, code a non-VIPA address for NSINTERADDR.
- If you use a name server to resolve host names and your name server runs on a z/OS Communications Server node that does not use VIPA addresses or if the name server runs on a non-z/OS Communications Server node, there are no additional coding requirements to TCPIP.DATA.

Step 3: Prepare your VTAM definitions

There are four categories of definitions you need for your APPN-enabled VTAM:

- Start options
- EE XCA major node
- Switched major node definitions
- Model major node definitions

Start options:

- EEVERIFY

The default value for this option is `ACTIVATE`, which is the recommended value for EE networks. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the EEVERIFY start option.

- HPRARB

The default value for this option is `RESPMODE`, which is the recommended value for EE networks. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the HPRARB start option.

- HPRPSDLY

Use this start option to specify the minimum amount of time that VTAM delays before an RTP pipe (that has an unresponsive partner) enters a path switch state. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the HPRPSDLY start option.

- HPRPSMSG

Indicates whether VTAM should limit the number of HPR path switch messages that are issued in the event of a network outage. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the HPRPSMSG start option.

- HPRPST

The maximum time that VTAM tries a path switch before ending a connection. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the HPRPST start option.

- HPRSESLM

Specifies whether VTAM should limit the number of sessions assigned to each RTP connection. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the HPRPST start option.

- HPRSTALL

Use this start option to set the amount of time an RTP pipe may remain continuously stalled before being stopped automatically. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the HPRSTALL start option.

- IPADDR/HOSTNAME

Use one of these options to specify a local static VIPA address that other EE endpoints will use to communicate with this VTAM. The decision whether to use IPADDR or HOSTNAME should be based on what you decided in the previous section, [“Step 2: Evaluate IP address resolution” on page 103](#).

- MULTPATH

Use this start option to disable IP multipath routing support for IPv4 and IPv6 Enterprise Extender connections. For more information about the MULTPATH start option, see [z/OS Communications Server: SNA Resource Definition Reference](#). For more information about the IP multipath routing, see [z/OS Communications Server: IP Configuration Guide](#).

- PMTUD

Use this start option to disable path MTU discovery for IPv4 and IPv6 Enterprise Extender connections. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the PMTUD start option.

- TCPNAME

See [TCPNAME](#) in [z/OS Communications Server: SNA Resource Definition Reference](#) for more information.

Tip: Although specifying TCPNAME is optional, you should specify a name for the TCP/IP stack that defines the local static VIPA for EE because it causes more reliable line activation. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the TCPNAME start option.

Restriction: All local Enterprise Extender static VIPA addresses must be associated with a single TCP/IP stack. Even in common-INET environments, EE always establishes affinity with only one TCP/IP stack.

- T1BUF/T2BUF

These are two buffer pool start options designed to optimize data transmission for Enterprise Extender configurations that use QDIO/iQDIO device drivers. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the T1BUF/T2BUF start options. Find out how tuning these options can benefit EE performance in [“Tuning the EE network” on page 130](#).

- UNRCHTIM

Specifies the options associated with Enterprise Extender connection network reachability awareness. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the UNRCHTIM start option start option.

EE XCA major node:

Activation of the EE XCA major node enables the use of EE, and establishes the following relationships.

- The PORT definition represents a relationship between VTAM and TCP/IP for sending and receiving EE packets.
- Each GROUP entry is composed of a set of LINE definitions sharing the same characteristics, such as IP address or host name. For groups that define a connection network, the APPN TG characteristics such as CAPACITY and SECURITY are defined on the GROUP statement in the XCA major node. The EEVERIFY operand on the GROUP statement overrides the EEVERIFY start option value. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the EEVERIFY operand under XCA major node for EE.

Guideline: Ensure that the CAPACITY specified is a value representing the effective capacity of the end-to-end IP connectivity. Note that this might be a different capacity than the TCP/IP interface.

- Each LINE definition represents a potential switched connection with a partner EE endpoint.

These relationships enable VTAM and TCP/IP to communicate across a connection that is associated with both the remote and local IP addresses. This communication begins when the first line, associated with a unique IP address, is activated.

The XCA major node defines the IP port with the following statements, operands and functions:

- The PORT statement identifies the name of a port through which an HPR connection through the IP network is made. The PORT statement must be set to MEDIUM=HPRIP for EE.

The IPPORT operand on the PORT statement specifies the first of five consecutive UDP ports intended for all EE communication from this host.

Requirements:

- The IPTOS operand on the PORT statement specifies values for the type of service bytes for each of the five UDP ports defined for EE. These IP precedence values within the ToS byte must be consistent throughout the network. By coding or defaulting the IPTOS operand, you preserve SNA transmission priority through the IP type of service byte that is recognized by IP networking hardware. Further customization of the ToS byte can be done through the z/OS IP policy agent. See the sections on IPTOS, XCA major node, Customizing IP type of service, and QoS preservation dependency in [z/OS Communications Server: SNA Resource Definition Reference](#) for more information.
- Use the default value of 12000 when specifying IPPORT.

The LIVTIME, SRQTIME, and SRQRETRY logical data link control (LDLC) parameters are specified on the PORT or GROUP statement. More information about these parameters is in [z/OS Communications Server: SNA Resource Definition Reference](#).

- The GROUP statement must be set to DIAL=YES to define a TG. Optionally, the GROUP statement can specify the source VIPA address explicitly or implicitly, using a host name value to be resolved into the address. If the IP address or host name is not defined on the XCA major node, the IP address or host name defined using VTAM start options is used as the source VIPA address. You must define separate GROUPs for IPv4 and IPv6 EE connections.

Guideline: Define at least as many lines as the maximum number of concurrent EE connections that are expected. Specifying DYNPU=YES on the GROUP statement indicates lines within the group that are eligible for dynamic PU assignments for inbound connections. This eliminates the need for explicit switched PU definitions for those connections. You can customize dynamic PUs for EE by defining a DYNTYPE=EE entry in a model major node. See [Model major node definition](#) for more information.

Requirements:

- For EE COS preservation, IP WAN devices must be enabled for QoS.
- All EE definitions must be coded within a single XCA major node.

Following is a sample XCA major node definition for Enterprise Extender:

```
*****
*
* NAME: XCA1A    XCA MAJOR NODE FOR HOST 1A Enterprise Extender
*
*****
XCA1A    VBUILD  TYPE=XCA
PORT1A   PORT    MEDIUM=HPRIP,IPPOR=12000,
              IPTOS=(20,40,80,C0,C0),LIVTIME=(10,30),
              SRQTIME=15,SRQRETRY=3

*
GP1A2A   GROUP   DIAL=YES,ANSWER=ON,ISTATUS=INACTIVE,
              CALL=INOUT,
              EEVERIFY=60,
              IPADDR=9.1.1.1,LIVTIME(15,40),SRQTIME=20,SRQRETRY=5

LN1A2A   LINE
P1A2A    PU

*
GP1A2A1  GROUP   DIAL=YES,ANSWER=ON,ISTATUS=INACTIVE,
              CALL=IN,DYNPU=YES,EEVERIFY=NEVER,
              IPADDR=9.1.1.2

LN1A2A2  LINE
P1A2A2   PU

*
* HOSTNAME resolves to IPv6 address
GP1A2A2  GROUP   DIAL=YES,ANSWER=ON,ISTATUS=INACTIVE,CALL=INOUT,
              HOSTNAME=HOST.DOMAIN.COM

LN1A2A2  LINE
P1A2A2   PU
*
* IPv6 address
```

```

GP1A2A3  GROUP  DIAL=YES,ANSWER=ON,ISTATUS=INACTIVE,CALL=INOUT,
              IPADDR=2001:0DB8::93:1:1:1
LN1A2A3  LINE
P1A2A3   PU

```

Switched major node definition:

The switched major node defines the Enterprise Extender switched TGs with the following statements:

- The PU statement defines the remote node
- The PATH statement defines the remote IPv4 address, or a host name that can be resolved into an IPv4 or an IPv6 address.

The PU definitions can be used for dialing outbound and for identifying PUs for incoming calls to this host. See the section on switched major node definitions in [z/OS Communications Server: SNA Resource Definition Reference](#) for more information.

Dial guidelines: For predefined EE connections:

- Coding a PATH definition statement with GRPNM enables an incoming call to prefer one switched PU over another. When the GROUP name that is associated with the line selected for the inbound call matches the GRPNM on the switched PATH, the switched PU associated with the PATH is preferred over another switched PU that does not match, or does not have a PATH definition statement coded.
- Specify APPN TG characteristics on the switched PU used for dial-out connections or on the DYNTYPE=EE model PU for dynamically defined dial-in connections.
- Specify the EEVERIFY operand on the switched PU used for dial-out connections or on the DYNTYPE=EE model PU for dynamically defined dial-in connections to override the EEVERIFY start option value. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about the EEVERIFY operand in the Switched major node or Model Major node sections, under the DYNTYPE=EE description.

Following is a sample switched major node definition for Enterprise Extender:

```

*****
*
*  NAME: SWXCA1A    SWITCHED MAJOR NODE FOR HOST 1A
*
*****
SWXCA1A  VBUILD  TYPE=SWNET,MAXNO=256,MAXGRP=256
*
SW1A2A   PU      IDBLK=003,IDNUM=00003,ADDR=01,
                  CPNAME=SSCP2A,EEVERIFY=60,
                  CPCP=YES,HPR=YES,
                  PUTYPE=2,TGP=GIGENET
PATH2A   PATH    IPADDR=93.1.1.1,
                  GRPNM=GP1A2A
*
*  HOSTNAME resolves to IPv4 address
SW1A2A1  PU      IDBLK=004,IDNUM=00004,ADDR=02,
                  CPNAME=SSCP2A,
                  CPCP=YES,HPR=YES,
                  PUTYPE=2,TGP=GIGENET
PATH2A1  PATH    HOSTNAME=SSCP2A,
                  GRPNM=GP1A2A1
*
*  HOSTNAME resolves to IPv6 address
SW1A2A2  PU      IDBLK=005,IDNUM=00005,ADDR=03,
                  CPNAME=SSCP2A,
                  CPCP=YES,HPR=YES,
                  PUTYPE=2,TGP=GIGENET
PATH2A2  PATH    HOSTNAME=SSCP2AV6,
                  GRPNM=GP1A2A2
*
*  IPv6 address
SW1A2A3  PU      IDBLK=006,IDNUM=00006,ADDR=04,
                  CPNAME=SSCP2A,
                  CPCP=YES,HPR=YES,
                  PUTYPE=2,TGP=GIGENET
PATH2A3  PATH    IPADDR=2001:0DB8::91:1:1:1,
                  GRPNM=GP1A2A3

```

Model major node definition:

To define the TG characteristics of a dynamic PU for EE, code a model EE PU within the model major node. This PU model definition is used only if all of the following are true:

- DYNPU=YES is explicitly coded on a GROUP statement under an EE PORT within an XCA major node
- A LINE statement coded under that GROUP statement is eligible to accept incoming calls (CALL=IN or CALL=INOUT, and ANSWER=ON are coded, sifted down from the GROUP statement, or taken as defaults)

Tip: Using CALL=INOUT provides the most flexibility.

- A predefined switched PU in connectable state that matches the partner cannot be found at dial-in time

Tip: Using the model PU enables you to specify things such as non-default TG characteristics and an installation-specific DISCNT value without having to code a PU statement for each remote node that can dial in. Also, by using the DWINOP, REDIAL, and REDDELAY parameters, you can specify whether you want VTAM to attempt to reconnect to the partner when the connection becomes inoperative. You can also specify up to four predefined TG numbers in the order you prefer for your EE connections. Include ANY as one of the four if you want to have the EE connection established even if your preferred TG numbers are not available. ANY can be placed anywhere in the list.

The EEVERIFY operand on the PU statement overrides the EEVERIFY start option value.

Following is a sample model major node definition that defines a model PU to be used for dynamically created, non-connection-network EE PUs for inbound connections.

```
*****
*
*   NAME: MODEL1A   MODEL MAJOR NODE FOR HOST 1A
*
*****
MODEL1A VBUILD   TYPE=MODEL
*
EEMODEL  PU      DYNTP=EE,
                  CAPACITY=100M,
                  COSTTIME=0,
                  CPCP=YES,
                  EEVERIFY=60,
                  DISCNT=NO,
                  DWINOP=YES,
                  REDIAL=30,
                  REDDELAY=60,
                  TGN=(11,8,15,ANY)
```

Activating EE

After you have done the preparation work, you can activate EE to receive benefits provided by EE.

Before you begin

Define static VIPAs to be used for EE connections, start the TCP/IP stack enabled for EE, and activate IUTSAMEH by static definition or by specifying DYNAMICXCF. See [“Step 1: Prepare your TCP/IP configuration”](#) on [page 102](#) for more information about preparing your TCP/IP configuration.

Procedure

Perform the following steps to start EE:

1. Enable z/OS Communications Server to accept or initiate EE connections by activating the EE XCA major node:
 - The first line activated in the EE XCA major node establishes a VTAM connection to the TCP/IP stack supporting Enterprise Extender. When this communication is established, the following messages appear on the console:

```
EZZ4313I INITIALIZATION COMPLETE FOR DEVICE IUTSAMEH
EZZ4324I CONNECTION TO 9.67.1.3 ACTIVE FOR DEVICE IUTSAMEH
```

- For the first line activation associated with each unique static VIPA address (IPADDR or HOSTNAME), VTAM registers this VIPA address with the TCP/IP stack for EE communications. Any other line you activate from a group results in a new message that indicates VIPA activation. The following message appears on the console:

```
EZZ4324I CONNECTION TO 9.67.1.5 ACTIVE FOR DEVICE IUTSAMEH
EZZ4324I CONNECTION TO 9.67.1.7 ACTIVE FOR DEVICE IUTSAMEH
```

Each message signifies that EE is registering the new static VIPA address with the TCP/IP stack and that the VIPA is ready to use.

2. Establish outbound EE connections by configuring a predefined EE network. Perform the following actions:

- Activate the EE switched major node.
- Dial the associated switched PU.

```
V NET,DIAL,ID=SWIP2A1

IST097I VARY ACCEPTED
IST590I CONNECTOUT ESTABLISHED FOR PU SWIP2A1 ON LINE LNIP1
IST1086I APPN CONNECTION FOR NETA.SSCP2A IS ACTIVE - TGN = 21
IST241I VARY DIAL COMMAND COMPLETE FOR SWIP2A1
IST1488I ACTIVATION OF RTP CNR00002 AS PASSIVE TO NETA.SSCP2A
IST1488I ACTIVATION OF RTP CNR00001 AS ACTIVE TO NETA.SSCP2A
IST1096I CP-CP SESSIONS WITH NETA.SSCP2A ACTIVATED
```

If the partner is ready to accept connections, connections will be established. To establish connections using EE connection network, see [“Configuring the EE connection network” on page 116](#).

3. Accept inbound EE connections by activating a switched PU to receive inbound connections.

Use one of the following methods to identify the switched PU that will receive the EE connection.

- Activate a switched PU to receive the incoming connection.
- Specify DYNPU=YES on a GROUP within the XCA major node to indicate that a dynamic PU should be created when inbound connections arrive. Optionally, specify a model to be used for creation of dynamic physical units for incoming calls by activating a model major node containing a PU statement with DYNTYPE=EE specified.

Results

Activating EE automatically:

If you want to automate EE activation and connection to EE partners with minimal operator intervention, specify the EE XCA major node and any related EE switched major nodes in the VTAM configuration start list (ATCCONxx). Specifying DWACT=YES on the switched PU causes VTAM to automatically dial the switched PU upon activation.

Tip: The connection attempt might initially fail. Consider coding REDIAL and REDDELAY on the switched path definition statement. For more information about path definition statements see [z/OS Communications Server: SNA Resource Definition Reference](#). For EE to automatically recover after a connection loss, consider coding DWINOP on the switched PU.

Guideline: Use caution if both EE endpoints support DWINOP. See [“Dial usability - DWACT, DWINOP, KEEPACT, REDIAL, and REDDELAY” on page 148](#) for more information.

Verifying activation

The following commands are useful in verifying that your EE network is correctly configured and performing properly.

- DISPLAY EE

This command displays general Enterprise Extender (EE) information.

```
d net,ee

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2000I ENTERPRISE EXTENDER GENERAL INFORMATION
IST1685I TCP/IP JOB NAME = TCPCS
IST2003I ENTERPRISE EXTENDER XCA MAJOR NODE NAME = XCAIP
IST2004I LIVTIME = (10,0)      SRQTIME = 15  SRQRETRY = 3
IST2005I IPRESOLV = 0
IST924I -----
IST2006I PORT PRIORITY = SIGNAL    NETWORK    HIGH    MEDIUM    LOW
IST2007I IPORT NUMBER = 12000      12001      12002      12003      12004
IST2008I IPTOS VALUE  = C0          C0          80          40          20
IST924I -----
IST2017I TOTAL RTP PIPES = 2        LU-LU SESSIONS = 2
IST2018I TOTAL ACTIVE PREDEFINED EE CONNECTIONS = 1
IST2019I TOTAL ACTIVE LOCAL VRN EE CONNECTIONS = 0
IST2020I TOTAL ACTIVE GLOBAL VRN EE CONNECTIONS = 0
IST2021I TOTAL ACTIVE EE CONNECTIONS = 1
IST314I END
```

For more information about detailed breakdowns of the DISPLAY outputs, see [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#).

- DISPLAY EEDIAG

This command displays Enterprise Extender (EE) connections that meet or exceed a specified retransmission threshold.

```
d net,eediag,rexmit=0,id=swip2a1

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2065I ENTERPRISE EXTENDER CONNECTION REXMIT INFORMATION
IST2067I EEDIAG DISPLAY ISSUED ON 06/13/05 AT 12:01:03
IST924I -----
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.2
IST1909I REMOTE HOSTNAME SSCP2A
IST2023I CONNECTED TO LINE LNIP1
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I NLPS SENT = 42799 ( 042K )
IST2037I BYTES SENT = 2995730 ( 002M )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 44186 ( 044K )
IST2041I BYTES RECEIVED = 3068884 ( 003M )
IST314I END
```

For more information about detailed breakdowns of the DISPLAY outputs, see [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#).

- DISPLAY TOPO

This command displays node and link information from the topology data base.

```
d net,topo,orig=dom,dest=chile,appncos=#inter

IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP NETNORTH.DOM
IST1357I
IST1300I DESTINATION CP      TGN      STATUS    TGTYPE    CPCP
IST1301I NETNORTH.CHILE    21      OPER      INTERM    VALUE WEIGHT
IST1579I -----
IST2241I
IST1163I RSN                HPR                TIME    ISL
IST1164I 172654            YES                15      WEIGHT
IST1579I -----
IST1302I CAPACITY PDELAY    COSTTIME    COSTBYTE
IST1303I 25M          NEGLIGIB 0              0
IST1579I -----
IST1304I SECURITY UPARM1    UPARM2      UPARM3
```

```

IST1305I          SECURE      128      128      128
IST1579I          -----
IST1736I          PU NAME
IST1737I          ISTPDOCH
IST924I  -----
IST2275I  TDU INFORMATION SINCE LAST RESET ON 01/23/14 AT 07:24:31
IST1769I  LAST TDU RECEIVED - 01/23/14 08:48:27 FROM NETNORTH.SWEDEN
IST2281I  LAST TDU SENT - 01/23/14 10:43:34
IST2282I  TDU COUNTS:
IST2352I      SENT      = 9              RECEIVED = 2
IST2353I      ACCEPTED = 0              REJECTED = 1
IST2354I      IGNORED  = 1
IST314I  END

```

For more information about DISPLAY TOPO see z/OS Communications Server: SNA Operation.

- DISPLAY ID=EE *puname*

This command displays information about the EE switched PU such as local and remote IP addresses.

```
d net,id=swip2a1
IST097I DISPLAY ACCEPTED
IST075I NAME = SWIP2A1, TYPE = PU_T2.1
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1043I CP NAME = SSCP2A, CP NETID = NETA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST1105I RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
IST1106I SWIP2A1 AC/R 21 YES 98750000000000000000000017100808080
IST1482I HPR = RTP - OVERRIDE = N/A - CONNECTION = YES
IST1510I LLERP = NOTPREF - RECEIVED = NOTALLOW
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.2
IST1909I REMOTE HOSTNAME SSCP2A
IST2114I LIVTIME: INITIAL = 10 MAXIMUM = 0 CURRENT = 10
IST136I SWITCHED SNA MAJOR NODE = TOIP2A
IST081I LINE NAME = LNIP1, LINE GROUP = GPIIP, MAJNOD = XCAIP
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST314I END
```

For more information about [DISPLAY ID command](#) in *z/OS Communications Server: SNA Operation* see *z/OS Communications Server: SNA Operation*.

- Verify the connections by issuing the appropriate VTAM DISPLAY command:

```
d net,ee,id=swip2a1
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST200I ENTERPRISE EXTENDER CONNECTION INFORMATION
IST075I NAME = SWIP2A1, TYPE = PU_T2.1
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.2
IST1909I REMOTE HOSTNAME SSCP2A
IST2346I CP NAME NETA.SSCP2A
IST2022I EE CONNECTION ACTIVATED ON 06/13/05 AT 11:53:51
IST2114I LIVTIME: INITIAL = 10 MAXIMUM = 0 CURRENT = 10
IST2023I CONNECTED TO LINE LNIP1
IST2025I LDLC SIGNALS RETRANSMITTED AT LEAST ONE TIME = 0
IST2026I LDLC SIGNALS RETRANSMITTED SQRETRY TIMES = 0
IST2009I RTP PIPES = 2 LU-LU SESSIONS = 2
IST2027I DWINOP = NO REDIAL = *NA* REDDELAY = *NA*
IST2028I KEEPACT = YES
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I NLPS SENT = 31 ( 000K )
IST2037I BYTES SENT = 2467 ( 002K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 32 ( 000K )
IST2041I BYTES RECEIVED = 2545 ( 002K )
IST314I END
```

or

```
d net,ee,cpname=sscp2a
IST097I DISPLAY ACCEPTED
```

```

IST350I DISPLAY TYPE = EE
IST2001I ENTERPRISE EXTENDER CONNECTION INFORMATION
IST924I -----
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.2
IST1909I REMOTE HOSTNAME SSCP2A
IST2346I CP NAME NETA.SSCP2A
IST2022I EE CONNECTION ACTIVATED ON 06/13/05 AT 11:53:51
IST2114I LIVTIME: INITIAL = 10 MAXIMUM = 0 CURRENT = 10
IST2023I CONNECTED TO LINE LNIP1
IST2024I CONNECTED TO SWITCHED PU SWIP2A1
IST2025I LDLC SIGNALS RETRANSMITTED AT LEAST ONE TIME = 0
IST2026I LDLC SIGNALS RETRANSMITTED SRQRETRY TIMES = 0
IST2009I RTP PIPES = 2 LU-LU SESSIONS = 2
IST2027I DWINOP = NO REDIAL = *NA* REDDELAY = *NA*
IST2028I KEEPACT = YES
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I NLPS SENT = 31 ( 000K )
IST2037I BYTES SENT = 2467 ( 002K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 32 ( 000K )
IST2041I BYTES RECEIVED = 2545 ( 002K )
IST2042I 1 OF 1 EE CONNECTIONS DISPLAYED
IST314I END

```

- **DISPLAY NETSTAT DEVLINKS**

This command generates a report that displays the information about devices and their interfaces or links that are defined to the TCP/IP stack. For more information about the DISPLAY NETSTAT command see [z/OS Communications Server: IP System Administrator's Commands](#).

```

d tcpip,,netstat,devlinks

.
.
DEVNAME: TRLE1A          DEVTYPE: MPCPTP
DEVSTATUS: READY
LNKNAME: TRLE1AL          LNKTYPE: MPCPTP      LNKSTATUS: READY
NETNUM: N/A  QUESIZE: N/A
ACTMTU: 14336
SECCLASS: 255
BSD ROUTING PARAMETERS:
MTU SIZE: N/A          METRIC: 00
DESTADDR: 9.67.16.2    SUBNETMASK: 255.0.0.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: YES
GROUP          REFCNT
-----
224.0.0.1      0000000001
LINK STATISTICS:
BYTESIN          = 12572401
INBOUND PACKETS  = 129102
INBOUND PACKETS IN ERROR = 0
INBOUND PACKETS DISCARDED = 0
INBOUND PACKETS WITH NO PROTOCOL = 0
BYTESOUT         = 12244450
OUTBOUND PACKETS = 125059
OUTBOUND PACKETS IN ERROR = 0
OUTBOUND PACKETS DISCARDED = 0
DEVNAME: IUTSAMEH      DEVTYPE: MPCPTP
DEVSTATUS: READY
LNKNAME: TOVTAM          LNKTYPE: MPCPTP      LNKSTATUS: READY
NETNUM: N/A  QUESIZE: N/A
ACTMTU: 65535
SECCLASS: 255
BSD ROUTING PARAMETERS:
MTU SIZE: N/A          METRIC: 00
DESTADDR: 0.0.0.0      SUBNETMASK: 255.0.0.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: YES
GROUP          REFCNT
-----
224.0.0.1      0000000001
LINK STATISTICS:
BYTESIN          = 0
INBOUND PACKETS  = 0
INBOUND PACKETS IN ERROR = 0

```

```

      INBOUND PACKETS DISCARDED          = 0
      INBOUND PACKETS WITH NO PROTOCOL  = 0
      BYTESOUT                          = 0
      OUTBOUND PACKETS                  = 0
      OUTBOUND PACKETS IN ERROR         = 0
      OUTBOUND PACKETS DISCARDED        = 0
      .
      .
      40 OF 40 RECORDS DISPLAYED

```

In addition to DISPLAY commands, z/OS Communications Server has a network management interface that provides EE configuration and performance data. There are various network management applications that use this interface to provide a graphical presentation of your EE network. An example is Tivoli ITM/NP V2. See [z/OS Communications Server: IP Programmer's Guide and Reference](#) for more information about the z/OS network management interface

Verifying the health of EE connection

VTAM sends a Logical Data Link Control (LDLC) probe to the remote partner to determine whether all ports are accessible during the activation of the EE connection by default. EE health verification for this connection activation fails if VTAM does not receive a response, or receives an error response, from the remote partner.

VTAM also verifies the health of EE connections at other user-specified time intervals. To do this, VTAM sends an LDLC probe to the remote partner on a user-specified interval. If the LDLC probe does not reach any port, VTAM issues the eventual action error message if it is not already present. The eventual action error message stays on console until the customer clears the message or all active EE connections receive successful EE health verification during the most recent LDLC probe.

The D NET,EE,LIST=EEVERIFY command displays all lines that failed EE health verification during the most recent LDLC probe.

The D NET,EE,ID=puname or linename command displays EE health verification information of the most recent LDLC probe for a particular connection.

VTAM provides the EE health verification information of the most recent LDLC probe to the NMI application upon its request. It includes success and failure information for all ports and contains the round-trip time information.

VTAM issues the health verification failure message if health verification fails during the activation of EE connection when it sends the LDLC probe to their remote partner.

```
IST2330I EE HEALTH VERIFICATION FAILED FOR puname AT time
```

During the activation of the EE connection, VTAM sends Logical Data Link Control (LDLC) probes to the remote partner to determine whether all the five ports are accessible. VTAM does not receive a response to any of the LDLC probe requests. VTAM continues with the activation of the EE connection between this node and the remote partner. Because VTAM receives no replies to its LDLC probe requests, VTAM determines that the remote partner does not support EE health verification and VTAM issues the following message:

```
IST2342I EE HEALTH VERIFICATION NOT SUPPORTED BY puname
```

When the EE Health verification fails on active connections, VTAM issues the following eventual action message. This message remains on the console until a subsequent LDLC probe to the remote partner is successful or is erased by the operator.

```
IST2323E EE HEALTH VERIFICATION FAILED FOR ONE OR MORE CONNECTIONS
```

When you see the above message, you can issue the D NET,EE,LIST=EEVERIFY command to display each line, which failed health verification on the most recent LDLC probe to its remote partner. The command displays the following message for each line with the line name, pu name, date and time.

```
IST2325I LINE linename PU puname ON date AT time
```


Display EE,LIST=EEVERIFY command example

This output shows all lines with active EE connections and failed EE health verification on last LDLC probe.

```
D NET,EE,LIST=EEVERIFY
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2000I ENTERPRISE EXTENDER GENERAL INFORMATION
IST1685I TCP/IP JOB NAME = TCPCS1
IST2003I ENTERPRISE EXTENDER XCA MAJOR NODE NAME = XCAEE
IST2004I LIVTIME = (10,0) SRQTIME = 15 SRQRETRY = 3
IST2005I IPRESOLV = 0
IST2231I CURRENT HPR CLOCK RATE = STANDARD
IST924I -----
IST2006I PORT PRIORITY = SIGNAL NETWORK HIGH MEDIUM LOW
IST2007I IPPORT NUMBER = 12000 12001 12002 12003 12004
IST2008I IPTOS VALUE = C0 C0 80 40 20
IST924I -----
IST2324I EE HEALTH VERIFICATION: FAILED CONNECTION INFORMATION
IST2325I LINE LNEE4000 PU SWEE42AI ON 08/11/09 AT 20:52:43
IST2326I EE HEALTH VERIFICATION TOTAL CONNECTION FAILURES = 1
IST2017I TOTAL RTP PIPES = 5 LU-LU SESSIONS = 5
IST2018I TOTAL ACTIVE PREDEFINED EE CONNECTIONS = 1
IST2019I TOTAL ACTIVE LOCAL VRN EE CONNECTIONS = 0
IST2020I TOTAL ACTIVE GLOBAL VRN EE CONNECTIONS = 0
IST2021I TOTAL ACTIVE EE CONNECTIONS = 1
IST314I END
```

This output shows EE health verification message from the most recent LDLC probe.

```
D NET,EE,ID=SWEE2A1
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2001I ENTERPRISE EXTENDER CONNECTION INFORMATION
IST075I NAME = SWEE2A1, TYPE = PU T2.1
IST1680I LOCAL IP ADDRESS FC00::67:1:1
IST1910I LOCAL HOSTNAME VIPA16.SSCP1A.TCP.RALEIGH.IBM.COM
IST1680I REMOTE IP ADDRESS FC00::67:1:2
IST1909I REMOTE HOSTNAME VIPA16.SSCP2A.TCP.RALEIGH.IBM.COM
IST2346I CP NAME NETA.SSCP2A
IST2114I LIVTIME: INITIAL = 10 MAXIMUM = 30 CURRENT = 20
IST2022I EE CONNECTION ACTIVATED ON 08/11/09 AT 20:22:55
IST2023I CONNECTED TO LINE LNEE1000
IST2327I EE HEALTH VERIFICATION OPTION - EEVERIFY = 3 MINUTES
IST2328I EE HEALTH VERIFICATION FAILED ON 8/11/09 AT 20:55:57
IST2339I EE HEALTH VERIFICATION LAST SUCCESS ON 08/11/09 AT 20:30:00
IST2025I LDLC SIGNALS RETRANSMITTED AT LEAST ONE TIME = 0
IST2026I LDLC SIGNALS RETRANSMITTED SRQRETRY TIMES = 0
IST2009I RTP PIPES = 4 LU-LU SESSIONS = 3
IST2027I DWINOP = NO REDIAL = *NA* REDDELAY = *NA*
IST2028I KEEPACT = NO
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I NLPS SENT = 49782 ( 049K )
IST2037I BYTES SENT = 4576487 ( 004M )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 49780 ( 049K )
IST2041I BYTES RECEIVED = 4576724 ( 004M )
IST314I END
```

Display EE,ID=line name command example

Display NET,EE,ID=line name shows EE health verification from the most recent LDLC probe.

```
D NET,EE,ID=LNEE1000,LIST=DETAIL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2001I ENTERPRISE EXTENDER CONNECTION INFORMATION
IST075I NAME = LNEE1000, TYPE = LINE
IST1680I LOCAL IP ADDRESS FC00::67:1:1
IST1910I LOCAL HOSTNAME VIPA16.SSCP1A.TCP.RALEIGH.IBM.COM
IST1680I REMOTE IP ADDRESS FC00::67:1:2
IST1909I REMOTE HOSTNAME VIPA16.SSCP2A.TCP.RALEIGH.IBM.COM
IST2346I CP NAME NETA.SSCP2A
IST2114I LIVTIME: INITIAL = 10 MAXIMUM = 30 CURRENT = 20
IST2022I EE CONNECTION ACTIVATED ON 08/11/09 AT 14:14:26
```

```

IST2024I CONNECTED TO SWITCHED PU SWEE2A1
IST2327I EE HEALTH VERIFICATION OPTION - EEVERIFY = 3 MINUTES
IST2328I EE HEALTH VERIFICATION FAILED ON 8/11/09 AT 20:55:57
IST2339I EE HEALTH VERIFICATION LAST SUCCESS ON 08/11/09 AT 20:30:00
IST2025I LDLC SIGNALS RETRANSMITTED AT LEAST ONE TIME      =      0
IST2026I LDLC SIGNALS RETRANSMITTED SRQRETRY TIMES        =      0
IST2009I RTP PIPES = 4 LU-LU SESSIONS = 3
IST2027I DWINOP = NO REDIAL = *NA* REDDELAY = *NA*
IST2028I KEEPACT = NO
IST924I -----
IST2030I PORT PRIORITY = SIGNAL
IST2029I MTU SIZE = 1232
IST2036I NLPS SENT = 4 ( 000K )
IST2037I BYTES SENT = 526 ( 000K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 4 ( 000K )
IST2041I BYTES RECEIVED = 577 ( 000K )
IST924I -----
IST2031I PORT PRIORITY = NETWORK
IST2029I MTU SIZE = 1232
IST2036I NLPS SENT = 91 ( 000K )
IST2037I BYTES SENT = 7036 ( 007K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 90 ( 000K )
IST2041I BYTES RECEIVED = 6693 ( 006K )
IST924I -----
IST2032I PORT PRIORITY = HIGH
IST2029I MTU SIZE = 1232
IST2036I NLPS SENT = 56390 ( 056K )
IST2037I BYTES SENT = 5184997 ( 005M )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 56394 ( 056K )
IST2041I BYTES RECEIVED = 5185656 ( 005M )
IST924I -----
IST2033I PORT PRIORITY = MEDIUM
IST2029I MTU SIZE = 1232
IST2036I NLPS SENT = 0 ( 000K )
IST2037I BYTES SENT = 0 ( 000K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 0 ( 000K )
IST2041I BYTES RECEIVED = 0 ( 000K )
IST924I -----
IST2034I PORT PRIORITY = LOW
IST2029I MTU SIZE = 1232
IST2036I NLPS SENT = 0 ( 000K )
IST2037I BYTES SENT = 0 ( 000K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 0 ( 000K )
IST2041I BYTES RECEIVED = 0 ( 000K )
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I NLPS SENT = 56485 ( 056K )
IST2037I BYTES SENT = 5192559 ( 005M )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 56488 ( 056K )
IST2041I BYTES RECEIVED = 5192926 ( 005M )
IST314I END

```

Configuring the EE connection network

Your EE network might benefit from having an EE connection network. A connection network is a representation of a shared access transport facility (SATF) which enables nodes that identify their connectivity to the SATF by a common virtual routing node (VRN) to communicate without having individually defined connections to one another. [Figure 37 on page 117](#) shows VRN connectivity. A connection network is also referred to as a virtual routing node (VRN).




Figure 37. VRN connectivity

Before you begin, you should understand the benefits and liabilities of an EE connection network. These are shown in the following figures.

Figure 38 on page 118, when an LU-LU session between resources on Communications Server for Linux and HOSTB is set up, the optimal route is directly through the network. However, if Communications Server for Linux and HOSTB are not directly defined to each other or are not connected to the same connection network, the indirect route is selected through HOSTA.




Figure 38. VTAM routing with an SATF

Optimal route calculation is achieved in one of two ways:

- Meshed connection definitions (define connections from every node to every other node)
- Connection network definition (define a connection to a connection network or virtual node)

Defining connections between each pair of nodes enables optional route calculation, but requires $(n*(n-1))$ definitions (where n is the number of nodes). In [Figure 39 on page 119](#), for example, the LU-LU session no longer traverses HOSTA.

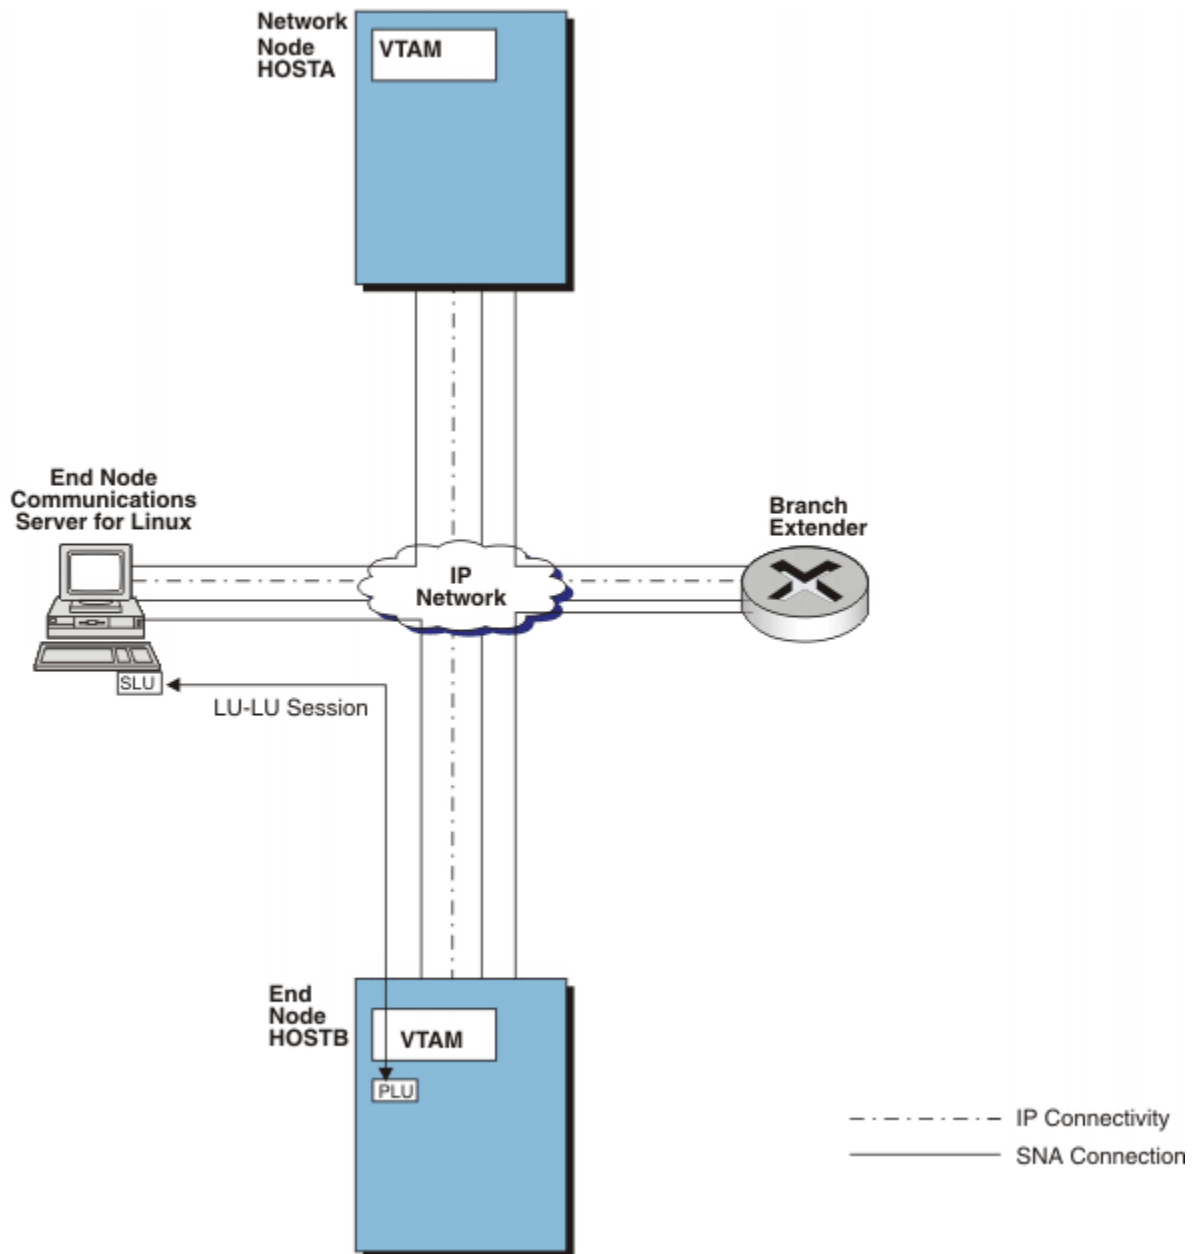


Figure 39. VTAM routing meshed connections

As the number of nodes in the network increases, the number of required connection definitions increases at an exponential rate. In a large network these definitions can be extensive. If you define a connection network to represent the shared access transport facility, system definition is reduced. With connection networks, end nodes need to define only a connection to a virtual node and to their network node server, which means only about $(2*n)$ connection definitions. [Figure 40 on page 120](#) shows this concept.




Figure 40. VTAM attachment to a connection network

The virtual node is reported to the topology database and can be chosen as the intermediate node during route calculation. This is shown in [Figure 41 on page 121](#).

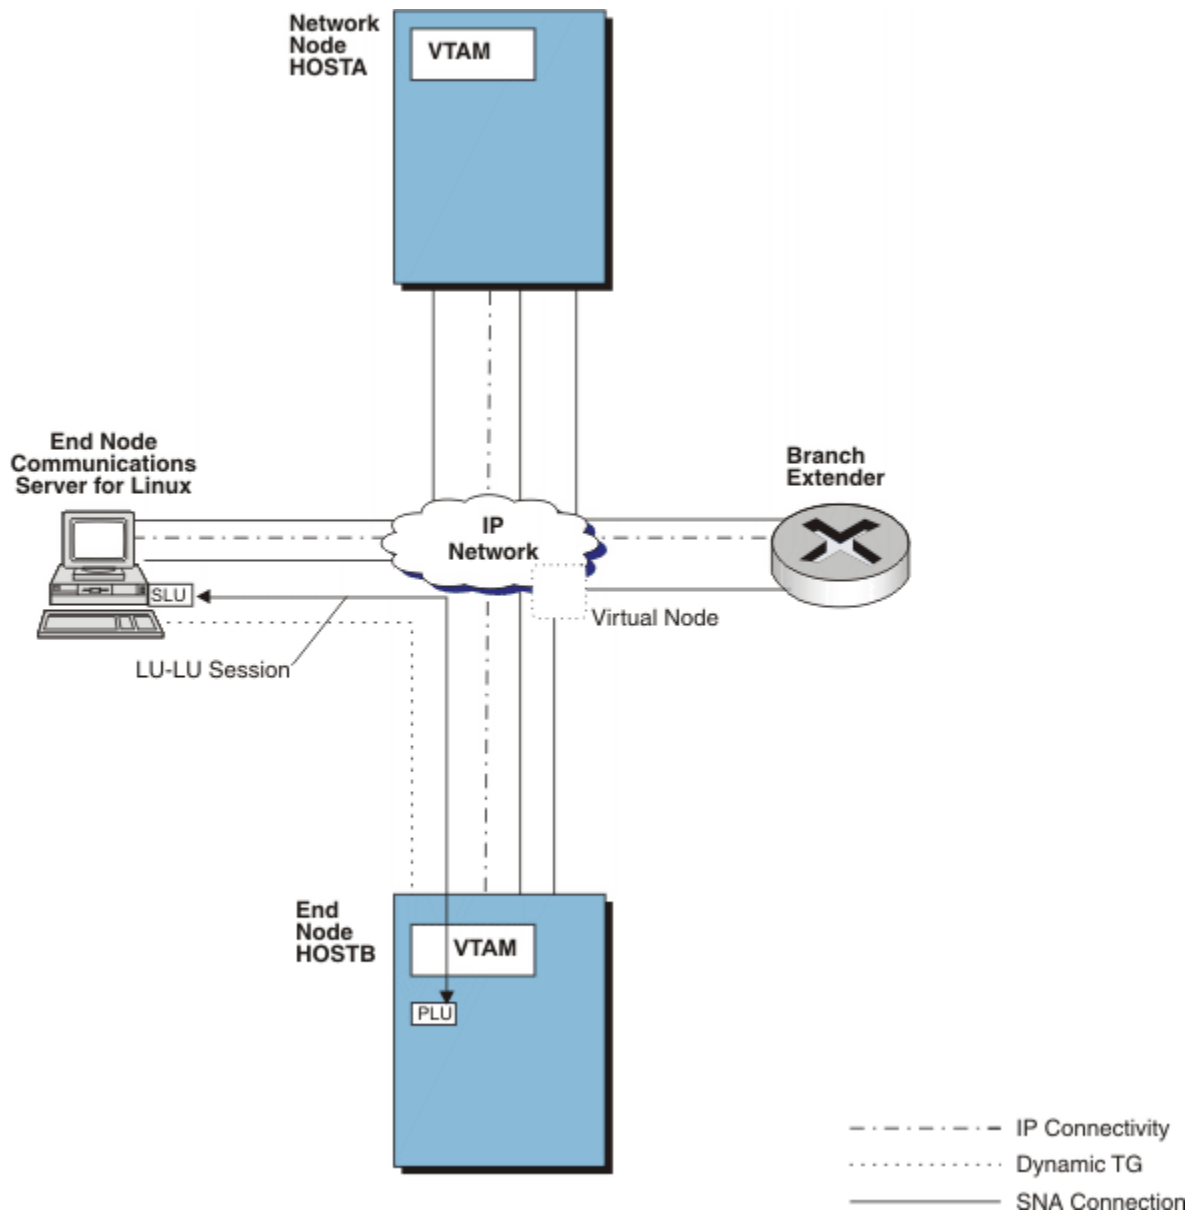


Figure 41. VTAM optimal route calculations

Connection network concepts

To understand connection networks, consider the following concepts:

- A connection network provides direct communication between all nodes on a shared access transport facility (such as an IP network) with the addition of just one definition on each node; namely, a link to the named connection network.
- A connection network is also referred to as a virtual node (VN) or a virtual routing node (VRN).
- Each EE connection network represents a single IP network of a given address family (IPv4 or IPv6) that is used for communicating between a set of EE-enabled partners.
- Connection network names are specified as fully qualified names of the form *netid.vnname*. All nodes on the same shared media that also have connection network functions can specify the same *netid.vnname*, which enables the node computing the route to recognize that all of these nodes are on the same shared media and can bringing up dynamic links for LU-LU sessions.
- Connection network links are not used for CP-CP sessions. End nodes need to predefine a link to their network node server (and usually to their backup network node server) and can also predefine links to other nodes. The predefined link to the network node server is used for CP-CP sessions. Dynamically

created connection network links are used only for LU-LU sessions and are typically brought down when the sessions end. (One exception to this is when a DYNTYPE=VN model is coded to change the setting of the DISCNT parameter for connection network PUs.) Similarly, network nodes also predefine links to adjacent network node partners for CP-CP sessions, with connection network links used only for LU-LU sessions.

- All participants on a local connection network (VNTYPE=LOCAL) must be located within the same APPN topology subnetwork. In a global connection network (VNTYPE=GLOBAL) nodes located in different APPN topology subnetworks can share the same connection network, thereby providing the capability for direct communication between the nodes defining the global VRN without requiring that all session traffic traverse the extended border nodes (EBNs) that interconnect those subnetworks. The EBNs must still be defined and participate in the session setup process, however.

EE connection network rules

- All partner nodes on an EE connection network must be able to communicate with each other over the IP WAN.
- Separate connection networks are needed for IPv4 and IPv6 traffic. If the TCP/IP stack supports both protocols, it can connect to both types of connection networks concurrently, if an IPv4 and an IPv6 local VIPA address have been defined to the stack.
- If host names are used for EE connection networks, then for consistent results, ensure that a given fully qualified host name resolves to a single source VIPA address. If not, the address returned by DNS becomes unpredictable.
- Host names are required when using IPv6 addressing. EE connection networks that use IPv6 addressing require name-to-address resolution for acquiring the source VIPA address of the local and remote EE endpoints.
- EE connection networks that use IPv4 addressing can also use name-to-address resolution for acquiring the source VIPA address of the local and remote EE endpoints, if both endpoints are z/OS V1R5 or later.

Contrasting local and global networks

If all connection network partners are located within the same APPN topology subnetwork, then you can define the connection network as a local connection network (VNTYPE=LOCAL). If some of the connection network partners are in different APPN topology subnetworks, then you must define the connection network as a global connection network (VNTYPE=GLOBAL).

Tip: Some products do not provide a choice between local and global connection networks. In general, if a product supports global connection networks but does not provide a choice of local or global (as part of connection network definition), then any connection networks that are defined are considered global.

Communication between partner nodes in different APPN topology subnetworks using a global connection network requires the assistance of extended border node connections. Based on these considerations, begin planning your EE connection network by drawing a logical diagram of your network showing each connection network and all participating partners.

Benefits of defining multiple Enterprise Extender virtual routing nodes

The following sections describe some of the benefits that can be achieved by defining multiple local VRNs (LVRNs), global VRNs (GVRNs), or both.

Defining local and global VRNs

You can define a local and global VRN to independently represent the connectivity within your own intranet, versus the connectivity over an extranet or the Internet. This enables you to provide direct access to specific systems for external vendors, while still giving internal systems direct access to all systems within your network. [Figure 42 on page 123](#) shows a network using multiple Enterprise Extender VRNs.

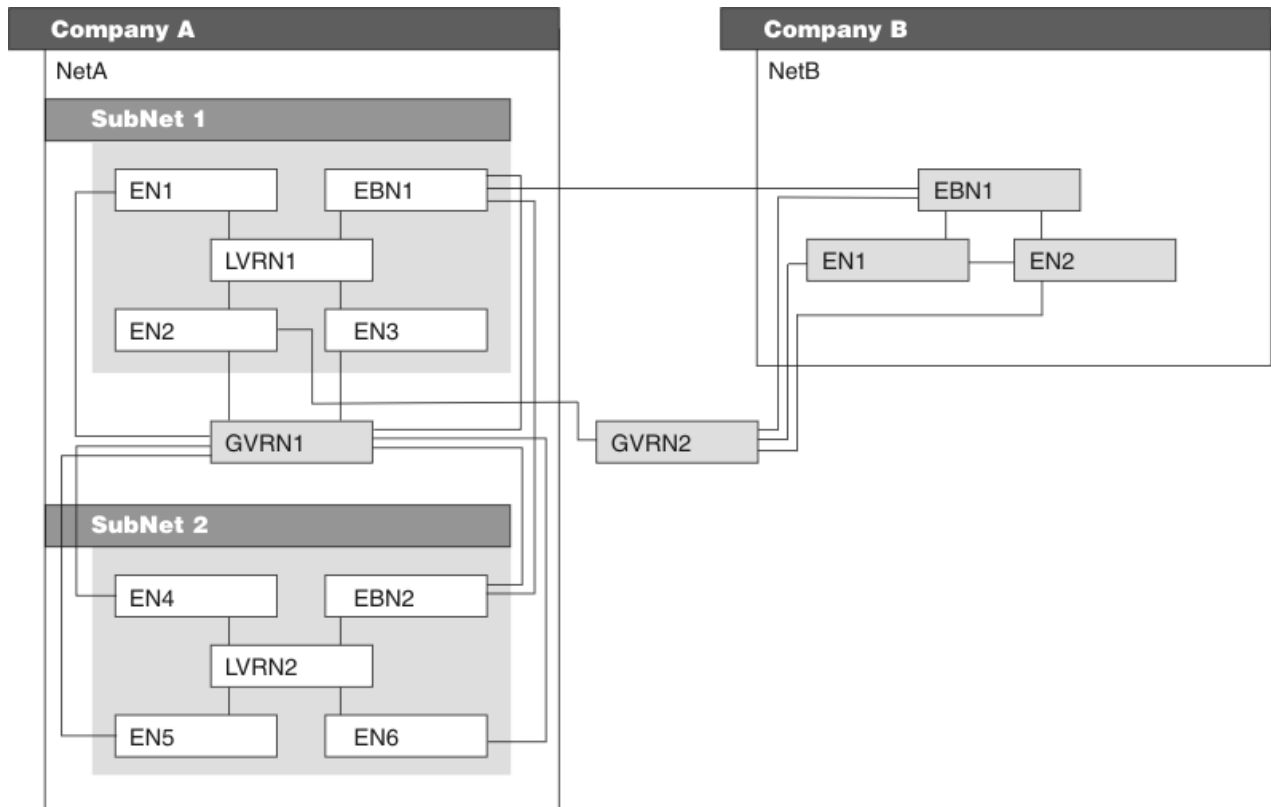


Figure 42. Defining multiple EE VRNs

In this figure, only the nodes in NetA SubNet 1 have access to the local VRN called LVRN1; and only the nodes in NetA SubNet 2 have access to LVRN2. These local VRNs can be used only to establish direct communication between nodes in the same subnetwork. But all of the nodes in both NetA SubNet 1 and NetA Subnet 2 have access to the global VRN called GVRN1. This allows every node in either of these two subnetworks to dynamically establish direct communication to any other node in either of these two subnetworks. GVRN2 is another global VRN which allows nodes in NetB to establish direct communication with only one end node (EN2) in NetA Subnet 1; sessions from NetB to any other node in NetA (SubNet 1 or SubNet 2) will include NetA.EBN1 as an intermediate node on the session path.

Defining multiple global VRNs

If a network is designed with multiple subnetworks, firewalls probably do not exist between these subnetworks. But if this same network also has APPN EBN connectivity to external vendors, a firewall probably exists between them. In this case, it would be beneficial to define two different global VRNs:

- A global VRN that is defined only by nodes within one of your own subnetworks
- A global VRN that is defined by your external vendor and a subset of the nodes within your own subnetworks

This enables you to control which systems external vendors can connect to directly (using global VRN), while still allowing internal systems (in any subnetwork) to directly connect to any of the other system in your subnetworks.

Additionally, defining multiple global VRNs enables you to avoid problems that might occur when multiple external customers connect to your network using the same global VRN. For example, assume Company A connects to Company B using a given global VRN (as shown in Figure 42 on page 123). If Company B connects to a second company using the same global VRN, it might then be possible for Company A to inadvertently connect directly to that second company. Although you can restrict this type of access by implementing the appropriate searching controls, defining a separate global VRN for connectivity to each external vendor increases your control over network access.

Using multiple local or global VRNs for EE load balancing

If you define links to multiple different VRNs (using different local IP addresses), you can choose different TG characteristics on these definitions to force sessions using different APPN COS names to flow over different VRNs (and use different local IP addresses). You can also configure the IP network to route traffic for different APPN COS names (IP ToS values) using different physical IP interfaces. For example, #INTER traffic could be sent over a Gigabit Ethernet connection and #BATCH traffic could be sent over a slower OSA link.

Using this type of configuration, you can monitor APPN traffic flows by displaying the RTPs that flow through a specific VRN. You can do this by specifying the `FIRSTCP=vrnname` operand, the `FIRSTTG=tgnumber` operand, or both on the `DISPLAY RTPS` command. This configuration also enables you to monitor IP traffic flows by tracing flows based on source or destination IP addresses.

If the APPN product you are using for a remote EE partner does not allow you to define multiple VRNs at the same time, then instead you can define two links to the same VRN at your z/OS system (using different local IP addresses) and force sessions to be load balanced over those connections. You can still monitor APPN and IP traffic by specifying both the `FIRSTCP=` and `FIRSTTG=` operands on the `DISPLAY RTPS` command in order to see what traffic is flowing over which link.

Finally, you might also want to use only a single local IP address, but still load balance sessions through parallel VRNs. With this type of configuration, it is not possible to monitor IP traffic over each of these connections, because there is only a single local IP address being used. But you would still be able to monitor APPN session traffic over the parallel routes (VRNs) by using the `FIRSTCP=` and `FIRSTTG=` operands on the `DISPLAY RTPS` command. Furthermore, because only a single IP address is being used, this type of configuration might simplify the coordination of system definitions.

Traversing multiple APPN network boundaries

To enable an EE connection network to traverse multiple APPN network boundaries, you can create an EE global connection network. In an EE global connection network, the following are true:

- An IP wide area network (WAN) is a shared access transport facility (SATF) represented by a global virtual routing node (GVRN).
- Nodes can indicate their connectivity to the SATF through the GVRN, thereby enabling nodes in the same or different APPN subnetworks to establish direct connectivity without having explicit routes (TGs) defined to one another.
- Part or all of a session path can traverse a common IP network (intranet or internet) and can bypass the extended border nodes (EBNs) completely.
- The computed session route does not have to traverse the adjacent border nodes (which is required when global VRNs are not used). This often provides for more optimal session routes and can remove the EBNs as a potential data bottleneck.

A GVRN can be used for cross-network EE connections if the following are true:

- The GVRN must be defined in the subnetwork of the proposed session endpoints.
- The GVRN must be defined on the endpoint node; or on the EBN that is acting as the entry point into (or exit point out of) the subnetwork of the endpoint node; or, under certain circumstances, on a network node in the subnetwork of the PLU. A GVRN defined on a network node can be used when that network node is in the same subnetwork as the PLU and the final session route is calculated by PLU node's network node server.
- The session path is completely APPN, and every subnetwork boundary along the session setup path is an extended subnetwork boundary. (That is, there is an EBN on both sides of every subnetwork boundary.) Each of the EBNs on the session setup path must be z/OS V1R2 Communications Server or later VTAMs.

Restrictions: The following restrictions apply when an EE global connection network is used for session establishment.

- GVRNs do not allow dynamic connections to end nodes residing below a branch extender node (BrNN). Assuming the BrNN has a connection to the GVRN, the BrNN is the concentration point for the EE global connection network for all nodes residing downstream of the BrNN.
- A GVRN is not used in intermediate subnetworks along a session setup path (except, possibly, as a local VRN).
- In most cases, a GVRN is not used if either session endpoint resides in or the session path traverses a subarea network (through an interchange node). However, the use of a DSME in intermediate border nodes might allow using a GVRN when the SLU resides in or the session path traverses a subarea network.
- GVRNs are not used for session setup requests that traverse APPN subnetwork boundaries defined with the value RTPONLY=YES. (The value RTPONLY=YES enables border nodes to maintain awareness of all sessions established over the specified APPN subnetwork boundaries; but using a GVRNs instead could result in an EBN not maintaining awareness of one or more sessions, which would defeat the purpose of RTPONLY=YES.) For more information about the RTPONLY operand, see [RTPONLY operand](#) in [z/OS Communications Server: SNA Resource Definition Reference](#).

Rules for identifying the source of the local IP address associated with a connection network:

- Each connection network partner must identify the IP address associated with its endpoint on the connection network. Specify either the host name or the IP address on the group definition where the connection network is defined.
- If the connection network represents an IPv6 network, then the connection network endpoint address must be specified using the HOSTNAME parameter.
- If the connection network represents an IPv4 network, then the endpoint can be defined by either the IP address parameter or the HOSTNAME parameter; however, the HOSTNAME parameter provides maximum flexibility because it provides compatibility with networks that use network address translation (NAT).

Defining an EE connection network in the EE XCA major node

Perform the following steps to configure the EE XCA major node to include the definition of a connection network.

1. Use the PORT and GROUP definition statements to define a connection to a connection network virtual routing node. Following is a sample Enterprise Extender XCA major node GROUP definition for a local VRN:

GRPEE1	GROUP	DIAL=YES,	(a)
		CALL=INOUT,	(b)
		HOSTNAME=NETA.HOST1,	(c)
		VNNAME=NETA.LCLVRN,	(d)
		VNTYPE=LOCAL,	(e)
		CAPACITY=1000M,	(f)
		AUTOGEN=(10,L,P)	(g)

Perform the following steps to define the appropriate operands:

- a. Code DIAL=YES on the GROUP statement that defines the virtual node or on the GROUP statement named on the VNGROUP operand of the PORT definition statement.
- b. Code the CALL operand to indicate that lines are available for both dial out and dial in connections.
- c. Code the HOSTNAME operand to specify the host name to be resolved to obtain this static VIPA address. Alternatively, code the IPADDR operand to specify the static VIPA address that Enterprise Extender connections will use to communicate with this host. Code the IPADDR or HOSTNAME on the GROUP definition statement, or it will default from the similarly named VTAM start option.

Tips:

- If your environment requires multiple source VIPA addresses, or if your EE connections require different addressing protocols (IPv4 versus IPv6), code the IPADDR operand or the HOSTNAME operand on the GROUP statement. The IPADDR operand explicitly identifies the source VIPA

address to be used for this EE connection. The HOSTNAME operand specifies the name to use for name-to-address resolution, at this node and at remote EE endpoints connected to this same VRN, for obtaining the source VIPA address representing this stack.

- If your environment requires unique LDLC inactivity timer settings for different EE connections, then you can specify different LIVTIME, SRQTIME, and SRQRETRY values on the GROUP definition statements (for each unique local IP address).
 - d. Code the VNNAME operand to define the name of the virtual node. IPv4 or IPv6 protocols can be used for both local or global VRN definitions. Supply a different virtual node name (VNNAME) for connection networks using different addressing protocols.
 - e. Code the VNTYPE operand to indicate whether the connection network is allowed to span APPN subnetwork boundaries. The default for VNTYPE is LOCAL (the connection network cannot span subnetwork boundaries). If the connection network should be allowed to span subnetwork boundaries, specify a VNTYPE of GLOBAL. (A VRN defined with VNTYPE=GLOBAL can also be used locally.) Attempts to define the same VRN as GLOBAL in some nodes and LOCAL in others might not produce the results you want. If a VRN is to be used across subnetwork boundaries, then define it as VNTYPE=GLOBAL in all nodes.
 - f. Use the CAPACITY operand to specify the media speed for the connectivity to the partner. The previous example specifies 1000 Mb for a gigabit Ethernet infrastructure. See [“Advanced coding considerations for EE”](#) on page 137 for more detail on specifying connection (TG) characteristics.
 - g. Coding AUTOGEN=(10,L,P) will automatically generate 10 lines and PUs for the group. The number of lines and PUs defined should be greater than or equal to the number of expected concurrent EE partners.
2. Define at least as many lines as the maximum expected concurrent connections. Because lines are associated with a specific GROUP statement, make sure that there are enough lines for that specific GROUP. For instance, a GROUP statement could be used for predefined EE connections (used for actual dialed EE connections) or for an EE connection network (either global or local). Make sure that there are enough EE lines for all GROUP statements defined in the XCA major node. These lines can be created manually or with AUTOGEN.
 3. Define as many local or global connection networks as your configuration requires. By allowing multiple global and local virtual routing nodes, users can define VRNs based on link characteristics. For example, a subset of users might require secure links, while others might use unsecured links. Depending on the requirements of the sessions, users can connect to the z/OS network using the appropriate VRN for the session characteristics. Definition of multiple local or global connection networks is necessary if both IPv4 and IPv6 protocols are to be used for local or global VRNs.
 4. To maximize availability, consider [“EE connection network reachability awareness”](#) on page 137.
 5. Activate the XCA major node if it is not already activated. Also activate the group by issuing VARY ACT, ID=GRPEE1.

You can verify that you have successfully defined a connection network in the EE XCA major node by looking for message:

```
IST1168I VIRTUAL NODE NETA.LCLVRN CONNECTION ACTIVE
```

which indicates that the connection network link is now active.

The following guidelines and restrictions apply.

Guidelines:

- When a VNTYPE of GLOBAL is specified and a VNNAME value is not specified, the virtual node is assigned the name IP.IP.
- The VNNAME/IPADDR pair (whether explicitly coded as an IP address using the IPADDR operand, or resolved from a host name using the HOSTNAME operand) must be unique. For example, if VRN1/HOST1 and VRN1/HOST2 are defined and if HOST1 and HOST2 resolve to the same IP address, then the first connection to be activated will succeed, but activation of the second connection will fail because these two definitions are not considered unique.

- You can define local connection networks (used for sessions within a single APPN topology subnetwork) and global connection networks (used for sessions that span APPN subnetwork boundaries). You can define as many local or global virtual nodes as your configuration requires. When a virtual node is defined on the PORT definition statement, the GROUP statement referenced by VNGROUP cannot also define a virtual node.

Restrictions:

- A local VRN cannot use the same name as a global VRN. A different VRN name (VNNAME) must be supplied for connection networks using different VNTYPES.
- If CP-CP sessions over EE are required between two nodes on the shared access transport facility, the dialing-out node must define the PU for every node that it will call. Routes traversing a virtual node cannot be used for CP-CP sessions.
- The dial-out node can use only a dynamically defined PU for connectivity through a connection network. The dial-in node attempts to use a predefined PU, if one exists. Otherwise, it uses a dynamically defined PU or a PU allocated using the configuration services XID exit.
- When a virtual node is defined on a GROUP definition statement, the TG characteristics for the link to that virtual node must also be specified on that GROUP definition statement.
- When a virtual node is defined on a PORT definition statement, the TG characteristics for the link to that virtual node must also be specified on the PORT definition statement.

EE security considerations

There are several techniques to consider when evaluating EE security options:

- SNA session level encryption (SLE)
- IPSec
- Network Address Translation (NAT)
- OEM security products

SNA session level encryption (SLE)

You can use SNA session-level encryption in the same way that it is currently used in SNA networks. In this way, the application can be assured that the traffic is received from the sender and that it has not been modified. This provides a reliable transmission method that is useful if sections of the network remain SNA (non-EE) and the data being transmitted is especially sensitive in nature. If the only section of the network that is unsecured is the IP network, or if authentication is the only requirement, other methods might work better for your needs.

Requirement: SLE requires the sharing of encryption keys. Ensure that you have a trusted partner.

IP security (IPSec)

IPSec is an industry-standard protocol that provides end-to-end authentication and encryption. IPSec provides an excellent method of securing EE connections. z/OS itself can be an IPSec endpoint or IP security can be offloaded to an attached router platform.

- By placing the IPSec endpoint on z/OS, you have end-to-end protection but your System z® CPU will incur the cost of the encryption.
 - Additional assistance for IPSec protocol traffic is available with any IBM Z Integrated Information Processor (IBM zIIP). You might need to modify some of your Workload Manager (WLM) definitions when you use zIIPs for IPSec-enabled EE connections. See [z/OS Communications Server: IP Configuration Guide](#) for information about modifying WLM definitions for zIIPs.
- By offloading the IPSec function to a router, you offload the encryption cost but you have an unprotected segment between z/OS and the router hosting the IPSec endpoint.

Providing authentication

If you need to provide authentication, you can use IPSec, which includes an authentication scheme. You can use IPSec authentication in one of the following ways:

- Stand-alone
- In combination with IPSec encryption
- In combination with SNA session-level encryption

Using IPSec's authentication scheme lets the system performing authentication verify that the data is received from the partner and has remained unchanged during transmission. When you use IPSec's authentication scheme in combination with encryption, you can authenticate the partner before wasting extra cycles for data decryption.

Tip: SNA also offers Message Authentication Code (MAC), but it does not authenticate the IP header portion.

Providing encryption

If the IP network is the only unsecured section, you can use IPSec between the two EE nodes to ensure that the transmitted data is not modified or viewed along the path.

- You can use IPSec between firewalls if there is a secure intranet and an unsecured Internet portion of the session path.
- You can run IPSec on the host to establish a VPN.
- If the EE nodes are the session partners, you can use either SNA session-level encryption or IPSec to encrypt the data.

The most significant difference between IPSec and SLE is that IPSec encrypts part of the UDP header, but SNA session-level encryption does not. See [z/OS Communications Server: SNA Resource Definition Reference](#) for specifics about session-level encryption.

Tip: The SNA header is encrypted only if IPSec is used.

If you use SNA encryption, use the filtering rule on the EE UDP port to allow traffic to flow without subsequent IPSec encryption. You can also use a combination of SNA encryption and IPSec authentication, where IPSec authentication is designed using filter rules on the same EE UDP port.

For more information about IPSec and IP filtering, see [z/OS Communications Server: IP Configuration Guide](#).

Using EE with network address translation (NAT)

Network address translation (NAT) is a technique where a one-to-one address translation function is performed, translating a single internal IP address to a single public IP address. NAT is a broad term that encompasses both a one-to-one address translation function, and network address/port translation (NAPT)

An internal-external IP address mapping is maintained by the NAT device. IP addresses are translated, but ports are unchanged. The mapping can be static or dynamic. For a static mapping, there is a definition in the NAT that always translates IP address x.x.x.x to IP address y.y.y.y. For a dynamic mapping, the NAT has a pool of IP addresses that are assigned as needed, so IP address x.x.x.x might be mapped to IP address y.y.y.y one time, and to IP address z.z.z.z at another time.

Restriction: Because EE requires a unique IP address mapping, dynamic network address translation is generally incompatible.

The following considerations apply when using EE on IP networks implementing network address translation:

- For predefined EE connections, ensure that the remote IP address is the public address translated or mapped by the NAT device.

- To enable EE connection networks to coexist with NAT, define the connection network using the HOSTNAME operand, not the IPADDR operand.

Network address/port translation (NAPT)

NAPT is a technique in which multiple internal IP addresses are translated into a single public IP address. As part of this translation process, the TCP and UDP ports in the packets are translated.

Restriction: Because EE has dependencies on port assignments, and a one-to-one IP address mapping, NAPT is incompatible with EE.

IP filtering

Use IP filtering to control the flow of network traffic. An IP security policy can define filters that deny or allow a packet access to a z/OS Communications Server system. Enterprise Extender (EE) IP filters use information in the IP and UDP packet headers to either permit or deny traffic coming into the z/OS CS enterprise. You can write an EE filter rule to allow inbound UDP data from a set of IP addresses to use the defined EE ports. You can write a second rule to allow outbound EE data access to a set of predefined IP addresses. You can also set up the policy to deny any traffic that does not match the previous rules. IP filtering rules can be written to log messages if a packet is denied because the packet did not match the rule. See [z/OS Communications Server: IP Configuration Guide](#) for more information about IP filtering.

IDS for Enterprise Extender

Use intrusion detection services (IDS) to protect z/OS Communications Server from attacks. You can use IDS to detect patterns that might indicate an impending attack. IDS supports attack rules that define the type of behavior to monitor. Define a set of attack rules to specify what events to monitor and the actions to take if an attack is detected. The action associated with the rule lets you specify what type of logging is to be performed and the action that the system should perform when an attack is detected.

Use the syslog daemon (syslogd) to log the attack. Syslogd is an application that writes messages to an OMVS file. You can set a statistics interval. At the end of each statistics interval, syslogd logs a message. The attack rule can write the detected packet to the IDS trace SYSTCPIS. The action can be defined to either discard or not discard the packet. You can code an exclusion list for each Enterprise Extender (EE) IDS attack type. Specific IP addresses in the exclusion list that would have been detected as an attack continue to be accepted to preserve existing behavior.

See [z/OS Communications Server: IP Configuration Guide](#) for more IDS information.

IDS defines the following new attack types for EE.

EE Malformed Packet

Specify this attack type to allow IDS to check inbound EE packets to determine whether the packet is malformed.

EE Port Check

Specify this attack type to allow IDS to verify the port value of inbound EE packets. If the port values are not correct, IDS flags the packet as a port check attack.

EE LDLC Check

Specify this attack type to allow IDS to verify EE LDLC commands. All EE data is sent by using LDLC commands. IDS checks the LDLC type to verify that the packet is received on the correct port.

EE XID Flood

Specify this attack type to allow IDS to detect suspicious XID activities and EE XID timeouts. EE XID timeouts might lead to an EE XID flood. When IDS detects an XID timeout, the following series of events occurs:

The local EE endpoint resends the XID reply three times before it fails the activation request and issues a timeout message. An XID flood occurs when the number of inbound XID timeouts is equal to the value of the EEXIDTimeout value in the IDS XID flood attack rule. Each inbound activation XID that is received by VTAM is assigned an available line for the connection. A partner EE that sends an XID and that does not continue activation of the connection will occupy an available line for about 1

minute. A flood of these occurrences can quickly use all available lines. This kind of attack is a denial of service attack. Valid XIDs will fail because a line is not available for the connection request.

The EEXIDTimeout value defines a threshold value that specifies the number of XIDs that can time out in 1 minute before IDS detects an EE XID flood. When IDS detects a flood, TCPIP writes a message to the console and to an OMVS file using syslogd (if such a message is required). The XID flood ends when the number of XIDs that are received in 1 minute is below the threshold value. You can enable statistics logging to assist in determining a threshold value. The statistics provide the number of XID timeouts that occurred during the interval and the maximum number of timeouts that occurred in any minute during the interval.

Note: The EE XID flood attack type does not support packet discard. It also does not support writing the packet to the IDS trace, SYSTCPIS.

OEM security products - EE proxy solutions

Because EE is based on UDP, firewalls in the underlying network must permit UDP traffic on the EE ports (12000 - 12004). For corporations with policies that absolutely forbid allowing UDP packets through the firewall, it is possible to proxy the UDP traffic onto TCP connections. Although IBM does not currently provide a TCP proxy for EE, other vendors might provide such a solution.

An additional advantage of using an EE proxy is that the EE traffic then becomes eligible for SSL encryption (SSL is a TCP protocol.) The disadvantage of an EE proxy (with or without SSL) is the additional CPU overhead incurred for the TCP transport.

Tuning the EE network

This section describes several ways to tune your EE network to improve performance.

Tuning Enterprise Extender-specific buffer pools

Enterprise Extender performance degradation can be caused by poorly tuned buffers. Monitor and tune the T1BUF and T2BUF buffer pools to minimize the number of expansions. Minimizing buffer pool expansions decreases internal buffer overhead processing, which should increase throughput while reducing CPU consumption. Also, the T1BUF and T2BUF pools are used instead of the TIBUF buffer pool. Increasing the buffer pool values for the T1BUF, T2BUF, or both might decrease the buffers necessary for the TIBUF pool. You should also monitor and tune the TIBUF pool.

Use IBM Health Checker for z/OS to check whether the T1BUF and T2BUF start option definitions for the initial number of buffers in the pool are greater than the default value when Enterprise Extender (EE) is not being used. You can also use Health Checker to check whether the initial number of buffers specified for these two buffer pools is equal to the default value when EE is being used with QDIO or HiperSockets. For more details about IBM Health Checker for z/OS, see [z/OS Communications Server: IP Diagnosis Guide](#).

The following is a buffer pool display example.

```
d net,bfruse,buf=(ti,t1,t2)
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = BUFFER POOL DATA
IST920I T100      BUFF SIZE  800      EXP INCREMENT  60
IST921I          TIMES EXP   0          EXP/CONT THRESH 120  / *NA*
IST922I          CURR TOTAL  360      CURR AVAILABLE  85
IST923I          MAX TOTAL  360      MAX USED        302
IST924I
-----
IST920I T100      BUFF SIZE  1340     EXP INCREMENT  36
IST921I          TIMES EXP   3          EXP/CONT THRESH 15   / 111
IST922I          CURR TOTAL  112      CURR AVAILABLE  112
IST923I          MAX TOTAL  112      MAX USED        69
IST924I -----
IST920I T200      BUFF SIZE  2028     EXP INCREMENT  32
IST921I          TIMES EXP 406920     EXP/CONT THRESH 7    / 103
IST922I          CURR TOTAL  104      CURR AVAILABLE  104
```


IST923I	MAX TOTAL	1608	MAX USED	1576
IST314I	END			

In this example, the T2BUF (T200) buffer pool expanded 406,920 times. At any given time, the maximum number of T2BUF buffers used at one time was 1576. This example indicates that the T2BUF buffer pool is thrashing, constantly expanding and contracting to meet the needs of its environment. A large number of expansions indicates that buffer pool tuning is warranted. In this case, coding a base allocation of 1576 buffers or more will eliminate the buffer pool thrashing for the T2BUF buffer pool. Use a similar approach when tuning your T1BUF buffers. In this example, the T1BUF pool does not need tuning (three expansions is considered minor).

Timers

EE and HPR depend on several different timers to determine when EE connections and RTP pipes should be terminated. There are timers for the RTP pipe itself, the EE switched PU, and at the logical data link control layer beneath the switched PU.

When does the RTP pipe go away?

With z/OS Communications Server, the RTP endpoint is represented by a dynamic PU, and acts like a delayed disconnect PU. This PU is dynamically given the name CNRxxxxx, where xxxxx represents a number incrementing from 00001 as the dynamic PUs are created. When the last session using the RTP terminates, the RTP goes inactive if no new session is queued to it before the disconnect time expires. You can alter the disconnect time using an operand on the DISCNT keyword on a model PU in a model major node (DYNTYPE=RTP indicates the model for HPR PUs). If you code a model RTP PU, code CPCP=YES if you want your CPCP sessions to traverse HPR pipes.

Guideline: Avoid using DISCNT=YES or a disconnect time of less than 10 seconds for DYNTYPE=RTP model PUs.

In the following example, 60 seconds is shown in the model major node definition for RTP PUs. Both RTP endpoints need to have similar types of definition. If dissimilar, the RTP endpoint with the lowest disconnect time would be the one deactivating the RTP, regardless of the (higher) value coded at the partner node.

MODMAJND	VBUILD	TYPE=MODEL
RTPPUMOD	PU	DYNTYPE=RTP, DISCNT=(DELAY, , 60), CPCP=YES

When does the EE connection go away?

Because the EE connection is itself represented by a switched PU, you can code the DISCNT operand to disconnect the EE switched PU. If DISCNT=YES is coded, the disconnect time used is that specified on the VTAM start option DISCNTIM (default is 15 seconds). Alternatively, you can specify DISCNT=(DELAY,,s) on an EE PU where s is equal to the number of seconds to be used for the disconnect time.

Because the EE PU (unlike RTP PUs) has no sessions queued directly to it, the disconnect decision is based on a period of inactivity rather than the termination of the last session.

If RTP's ALIVE timer is set to half of the disconnect time, RTP status messages (keep-alive flows) will keep the PU from dropping as long as at least one RTP is using the EE PU.

The Logical Data Link Control (LDLC) layer monitors the EE connection, and terminates the EE connection if contact is lost with the partner. The LDLC inactivity trigger is controlled by three parameters on the PORT or GROUP statements:

- LIVTIME - The amount of inactivity time that can lapse before LDLC tests the connection
- SRQTIME - The amount of time LDLC waits for a response to its test
- SRQRETRY - The number of times the test is retried

The connection is terminated if no activity or response occurs when using the following formula:

$LIVTIME + ((SRQRETRY+1) * SRQTIME)$ seconds

This is an approximate formula. The actual duration of the LDLC layer outage detection window should be within plus or minus one LIVTIME interval of the calculated amount.

Guidelines:

- Lengthen the HPR path switch timers (HPRPST) as necessary to ensure that all four timers are longer than the LDLC timeout interval. This ensures that RTP pipes stay in path switch long enough during IP network instability for the EE link to disconnect, and enables another path to be selected.
 - Alternatively, specify HPRPSDLY=EEDelay on the definition statements that represent EE connections. 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](#).
- When possible, ensure that the values of the LDLC parameters are consistent between the endpoints of the connection. If the EE partner is not a z/OS endpoint, see that product's documentation for details on setting its LDLC parameters.
- For predefined EE PUs, specify DISCNT=NO.

Tip: The LIVTIME, SRQTIME, and SRQRETRY parameters are customer-configurable parameters that control the ability of the LDLC-layer to detect TCP/IP network connectivity problems; these parameters should not be confused with the similarly-named variables associated with HPR's Adaptive Rate-Based congestion control (ARB) algorithm that is used to monitor each RTP connection. Each RTP has associated with it a liveness timer, a short request timer, and a short request retry counter. However, the RTP values are algorithmically maintained, are not configurable, and are independent of the LDLC parameters. (It is possible to influence the RTP's liveness timer value by coding DISCNT=YES for the connections used by the RTP.)

See the SRQRETRY description in topic in [z/OS Communications Server: SNA Resource Definition Reference](#) for more information.

Inactivity timer example 1

Consider the following time-annotated console:

```
d net,rtps
...
IST1695I PU NAME          CP NAME          COS NAME SWITCH CONGEST  SESSIONS
IST1696I CNR00004 NETA.SSCP2A #INTER    NO      NO      1
IST1696I CNR00003 NETA.SSCP2A RSETUP    NO      NO      0
...
(11:43:32) Break the IP connectivity.
v tcpip,,stop,trle1a
EZZ0060I PROCESSING COMMAND: VARY TCPIP,,STOP,TRLE1A
EZZ0053I COMMAND VARY STOP COMPLETED SUCCESSFULLY
EZZ4315I DEACTIVATION COMPLETE FOR DEVICE TRLE1A
(11:44:14) RTP CNR00004 detects failure and goes into path switch.
IST1494I PATH SWITCH STARTED FOR RTP CNR00004
(11:44:49) The LDLC layer detects loss of IP connectivity and
disconnects the EE connection.
IST1411I INOP GENERATED FOR LNIP1
IST1430I REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT
IST314I END
IST259I INOP RECEIVED FOR SWIP2A1 CODE = 01
IST619I ID = SWIP2A1 FAILED - RECOVERY IN PROGRESS
IST1196I APPN CONNECTION FOR NETA.SSCP2A INACTIVE - TGN = 21
IST590I CONNECTION TERMINATED FOR PU SWIP2A1 ON LINE LNIP1
IST621I RECOVERY SUCCESSFUL FOR NETWORK RESOURCE SWIP2A1
IST1488I INACTIVATION FOR RTP CNR00003 AS PASSIVE PARTNER COMPLETED
IST619I ID = CNR00003 FAILED - RECOVERY IN PROGRESS
IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00003 - VARY INACT SCHED
IST105I CNR00003 NODE NOW INACTIVE
IST871I RESOURCE CNR00003 DELETED
(11:46:14) CNR00004's path switch timer (2 min for interactive TP) expires.
CNR00004 goes inactive.
IST1494I PATH SWITCH FAILED FOR RTP CNR00004
IST1495I NO ALTERNATE ROUTE AVAILABLE
IST314I END
IST1488I INACTIVATION FOR RTP CNR00004 AS ACTIVE PARTNER COMPLETED
IST619I ID = CNR00004 FAILED - RECOVERY IN PROGRESS
```

```
IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00004 - VARY INACT SCHED
IST105I CNR00004 NODE NOW INACTIVE
IST871I RESOURCE CNR00004 DELETED
```

This is an example of how the various parameters influence the termination of an HPR connection using an EE TG. In this example, the RTP's disconnect time is the VTAM start option's DISCNTIM default, 15 seconds. The keyword DISCNT=YES is coded on the Enterprise Extender PU in its switched major node and, in the XCA major node on the PORT statement, LIVTIME is set to 10, SRQTIME to 15, and SRQRETRY to 3.

Based on the LDLC parameters, the LDLC layer should detect an outage and disconnect the connection in 70 seconds, plus or minus one SRQTIME interval. This results from the following equation:

$$10 + ((3+1) * 15) = 70 \text{ seconds}$$

With DISCNTIM set to 15 seconds, RTP connections using the EE connection will have a liveness timer value of 7.5 seconds (half the DISCNTIM value), which means that they will be likely to detect the outage faster than the LDLC layer.

As shown in the example, at 11:43:32 the IP TRLE is taken down and after about 40 seconds the RTP connection has detected the failure and has started a path switch. The LDLC layer detects the failure at 11:44:49, 77 seconds (close to the predicted value of 70 seconds) after the outage, and ends the EE connection. With no alternate path available, all path switch attempts fail, and the RTP pipe (CNR00004) fails at 11:46:14 (2 minutes after it went into path switch, based on the path switch timer value for interactive transmission priorities as set by HPRPST).

Inactivity timer example 2

Consider the following time-annotated console:

```
d net,rtps
...
IST1695I PU NAME          CP NAME          COS NAME SWITCH CONGEST  SESSIONS
IST1696I CNR00004 NETA.SSCP2A #INTER    NO      NO      1
IST1696I CNR00003 NETA.SSCP2A RSETUP    NO      NO      0
...
(11:34:56) Break the IP connectivity.
v tcpip,,stop,trle1a
EZZ0060I PROCESSING COMMAND: VARY TCPIP,,STOP,TRLE1A
EZZ0053I COMMAND VARY STOP COMPLETED SUCCESSFULLY
EZZ4315I DEACTIVATION COMPLETE FOR DEVICE TRLE1A
(11:36:23) The LDLC layer detects loss of IP connectivity and
           disconnects EE connection.
IST1411I INOP GENERATED FOR LNIP1
IST1430I REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT
IST314I END
IST259I INOP RECEIVED FOR SWIP2A1 CODE = 01
IST619I ID = SWIP2A1 FAILED - RECOVERY IN PROGRESS
IST1196I APPN CONNECTION FOR NETA.SSCP2A INACTIVE - TGN = 21
IST590I CONNECTION TERMINATED FOR PU SWIP2A1 ON LINE LNIP1
IST1494I PATH SWITCH STARTED FOR RTP CNR00004
IST621I RECOVERY SUCCESSFUL FOR NETWORK RESOURCE SWIP2A1
IST1488I INACTIVATION FOR RTP CNR00003 AS PASSIVE PARTNER COMPLETED
IST619I ID = CNR00003 FAILED - RECOVERY IN PROGRESS
IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00003 - VARY INACT SCHED
IST105I CNR00003 NODE NOW INACTIVE
IST871I RESOURCE CNR00003 DELETED
(11:38:23) CNR00004's path switch timer (2 min for interactive TP) expires.
           CNR00004 goes inactive.
IST1494I PATH SWITCH FAILED FOR RTP CNR00004
IST1495I NO ALTERNATE ROUTE AVAILABLE
IST314I END
IST1488I INACTIVATION FOR RTP CNR00004 AS ACTIVE PARTNER COMPLETED
IST619I ID = CNR00004 FAILED - RECOVERY IN PROGRESS
IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00004 - VARY INACT SCHED
IST105I CNR00004 NODE NOW INACTIVE
IST871I RESOURCE CNR00004 DELETED
```

In this example, the LDLC parameters (LIVTIME, SRQTIME, and SRQRETRY) yield a longer disconnect detection time of 95 seconds, as shown in the following equation. The keyword DISCNT=YES is coded on

the Enterprise Extender PU in its switched major node and, in the XCA major node on the PORT statement, LIVTIME is set to 20, SRQTIME to 15, and SRQRETRY to 4.

$$20 + ((4+1) * 15) = 95 \text{ seconds}$$

DISCNTIM is also much longer (240 seconds) than in the previous example, giving an RTP liveness timer value of 120 seconds.

In this example, the LDLC layer detects the outage first, generating an inoperative indication 87 seconds (again, close to the predicted value) after the IP connectivity is broken. In this case, the local link's inoperative transition forces the RTP connection CNR00004 to go immediately into path switch as can be seen by the IST1494I message several lines after the inoperative indication occurs at 11:36:23. Two minutes later (again based on the HPRPST value for interactive TP), at 11:38:23, the path switch timer expires and the RTP connection is terminated.

HPR ALIVE timer optimization for Enterprise Extender

In addition to the overhead of the timer maintenance itself, the I/O cost of sending keep-alive status messages and the processing of the subsequent replies has a significant impact on CPU utilization when a large number (thousands) of RTP pipes are active. You can use the HPREELIV start option to reduce the CPU impact.

The EE LDLC layer monitors the health of the TCP/IP network underlying the EE connection by monitoring for activity and sending test frames to verify partner reachability. In the event that the partner can no longer be reached, the LDLC layer generates a disconnection for the EE logical connection. For products supporting the architectural option of local link disconnection notifications to drive HPR path switch processing, this means that RTP pipes using EE connections will immediately enter path switch processing if the EE LDLC layer triggers an EE disconnection. Therefore, for RTP pipes that are composed of single EE connections (and including those pipes where the RSCV describes a two-hop path across an EE connection network), HPR ALIVE timer processing is superfluous. As long as the EE LDLC layer has not detected a problem, the RTP partner is considered reachable.

If HPREELIV is set to or defaults to YES, then:

- An RTP pipe with a route composed of a single EE link (including a single hop across a two-hop EE VRN path) can run without an ALIVE timer, thereby relying on the EE LDLC layer for inactivity monitoring, and saving the timer maintenance and keep-alive status request and reply overhead.
- RTP pipes composed of multi-hop paths, or single non-EE hops use ALIVE timer processing as architected.
- An RTP that path switches between one of these types of paths (for example, from single-hop EE to multi-hop) has to dynamically enable and disable the ALIVE timer at the time of the path switch.

There is CPU savings for configurations using HPR-over-one hop Enterprise Extender connections, or two hop Enterprise Extender virtual routing node (VRN) connections. When HPR-over-EE configurations specify DISCNT=NO, significant CPU savings can be achieved. Specifying DISCNT=NO means that HPR and EE connections always stay up, even when there are no sessions.

Enterprise Extender LDLC keep-alive reduction

During periods of relatively low network activity, configurations with many EE connections can consume a substantial amount of CPU resources because of the EE LDLC keep-alive function. In addition to the overhead of the timer maintenance itself, the I/O cost of sending keepalive messages (TEST request) and the processing of the subsequent responses has a significant impact on CPU utilization when large numbers of EE connections are active.

The Enterprise Extender LDLC keep-alive reduction function assists in reducing the total CPU impact of Enterprise Extender LDLC keep-alive processing overhead. You can specify an initial and a maximum value (through the max_value parameter) for the logical data link control (LDLC) liveness timer interval, using the LIVTIME operand (see LIVTIME operand in [z/OS Communications Server: SNA Resource Definition Reference](#) for more information). In this way, each Enterprise Extender connection can maintain

its own liveness timer value depending on network conditions. The liveness timer specifies how long the LDLC layer will wait without receipt of a network layer packet (NLP) from the connection partner before sending a liveness message (TEST frame) to verify that the connection is still operational. During periods of inactivity, an EE connection can expand its current liveness timer up to the limit specified by the `max_value` parameter. During each liveness interval, if no HPR data is received from the EE partner, the current liveness interval doubles in size. This process is repeated until the maximum liveness timer value is reached, or HPR traffic resumes. If the maximum liveness timer value is reached, the liveness timer will remain at that value. When HPR traffic resumes, the current liveness window is reset to the initial value and the EE connection is immediately tested to verify partner reachability.

The amount of CPU savings realized by LDLC keep-alive reduction is system-dependent. When HPR over EE configurations specify `DISCNT=NO`, significant CPU savings are possible.

Enterprise Extender improved packet loss tolerance

While RTP pipes generally perform well over Enterprise Extender under ideal network conditions, the overall performance of RTP pipes can be affected when packet loss is occurring. The Enterprise Extender improved packet loss tolerance function is designed to improve the performance of RTP pipes while running in non-ideal network conditions. This function is enabled by coding the `HPRCLKRT` start option to the `ADAPTIVE` mode.

Coding `ADAPTIVE` for the `HPRCLKRT` setting allows the HPR clock to operate in adaptive mode. Running the HPR clock in adaptive mode allows the clock to toggle between standard and high modes of operation when servicing RTP pipes which are experiencing packet loss. When packet loss occurs, the clock speed adapts to run in high mode. This causes more frequent scheduling of the RTP pipes which are experiencing packet loss, allowing the RTP pipes to recover quicker. There may be some CPU overhead for running the HPR clock in high mode. When packet loss is no longer occurring, the clock returns to the standard mode.

The adaptive mode is designed to increase the overall performance of RTP pipes when packet loss is occurring. However, RTP performance and CPU overhead are system-dependent and may vary accordingly.

Restrictions: When coding `HPRCLKRT=ADAPTIVE`:

- This mode is only valid in Enterprise Extender configurations which have a defined capacity of 1G (gigabit) or higher access speeds.
- This mode is not recommended for uni-processor configurations or in configurations which are CPU bound.

Disconnect and inactivity summary

The LDLC layer monitors the EE connection, sending a test frame if no activity is detected for the number of seconds specified by the `LIVTIME` operand as coded (or defaulted) on the `XCA PORT` statement. If no response is received for the number of seconds specified by the `SRQTIME` operand, another test frame will be issued. As long as no response is received, LDLC will retry `SRQRETRY + 1` times. If no response is received after LDLC tries again the last time, the EE link will be disconnected.

At a slightly higher layer, z/OS Communications Server monitors the EE switched PU if `DISCNT=YES` or `=DELAY` is specified.

- If you specify `YES`, a disconnect is generated if no traffic flows over that PU for the amount of time specified by the `DISCNTIM` start option.
- If you specify `DELAY`, a disconnect is generated if no traffic flows over that PU for the amount of time specified by the value you specified for the delay.

z/OS Communications Server provides the EE switched PUs disconnect time value to each RTP, and the RTP sets its liveness timer (which is independent from the LDLC liveness timer associated with the `LIVTIME` parameter) to half of the disconnect value, or 3 minutes, whichever is smaller. This ensures that as long as an RTP connection is actively associated with that EE switched PU, it sends status messages

(generated when the liveness timer pops) often enough to keep the PU from disconnecting because of inactivity, even if there is a lull in data traffic using the RTP connection.

At yet a higher layer, the RTP connection is itself represented by a switched PU that, by default, looks like a PU with DISCNT=YES and a 10 second DISCNTIM value. Therefore, by default the RTP PU will disconnect itself if its last session goes away, and no new session is queued to it for a period of 10 seconds. If a different DISCNTIM value is required for the RTP PU, or if DISCNT=NO behavior is required (meaning the RTP PU will not disconnect until forced down by lack of connectivity, partner deactivation, or operator command), then you can code a model PU, DYNTYPE=RTP.

Customizing IP type of service

EE maps the SNA transmission priority to the IP type of service so that routers can preserve the SNA transmission priority across the IP network.

The IP type of service (ToS) used for IPv4, and the IP traffic class used for IPv6, are related to the transmission priority of the chosen APPN class of service (COS) and the UDP port numbers used for EE traffic.

Tip: For IPv6 protocols, IP traffic class performs the same function as IP ToS does for IPv4 protocols. The term ToS is used in this document to see both functions.

The APPN COS specifies a transmission priority (with the keyword PRIORITY), which can be one of the data values shown in [Table 6 on page 136](#).

Restriction: Although the IPPORT operand enables you to change the port numbers used for EE, many platforms do not permit this. Therefore, do not change the IPPORT default value.

The UDP port value specified on the IPPORT keyword (default 12000), is reserved for signal data that is involved in establishing and keeping the EE connection active. The next four sequential UDP port numbers are reserved for the SNA traffic priorities (network, high, medium, and low).

IP routers can use the UDP port number to prioritize traffic. Likewise, by default, the ToS values in Table 5 tell IP routers that support this function how to prioritize data if they do not support prioritization by the port number. This also applies in cases where IPSec is being used.

Table 6. SNA priorities and corresponding port numbers and default ToS values		
Priority	Port	ToS value
LDLC Signaling *	12000	C0
Network	12001	C0
High	12002	80
Medium	12003	40
Low	12004	20

Note: * Port 12000 carries EE LDLC signaling data and does not correspond to one of the four SNA transmission priorities.

Not all routers support prioritization of traffic based on port number or ToS value. Some might support none and others might support one or both. This is why two methods are needed (both ToS and port number) for the prioritization of HPR traffic in an IP network. In addition, if IPSec is being used, EE data can be encrypted in the IP network. This encryption includes the UDP port numbers, but not the ToS values.

Guideline: For optimum preservation of SNA transmission priority, use prioritization based on ToS value.

Rule: If HPR data prioritization in an IP network is needed while using IPSec, the IP routers in the network must support the ToS function. Also, all routers and EE endpoints must be consistently defined for ToS and the UDP port numbers.

Advanced coding considerations for EE

The topics described in this section include:

- [“EE connection network reachability awareness ” on page 137](#)
- [“TCP/IP MTU size for EE ” on page 144](#)
- [“Running EE in constrained or virtualized environments” on page 145](#)
- [“RTP transmission stall operator awareness and recovery support” on page 146](#)
- [“Load balancing” on page 147](#)
- [“Transmission group profiles” on page 147](#)
- [“Dynamic reconfiguration” on page 148](#)
- [“Dial usability - DWACT, DWINOP, KEEPACT, REDIAL, and REDDELAY” on page 148](#)
- [“Customization for EE connection network PUs” on page 149](#)
- [“Cross-subnet routing with global VRNs” on page 150](#)

EE connection network reachability awareness

When a dial failure or a connection failure occurs for a connection over an Enterprise Extender connection network, the partner node cannot be reached using the connection network path across the virtual routing node (VRN). This connection network path might have been chosen for this connection because it had a lower weight than any alternate path available at one of the following times:

- The time of this failing dial
- The time of the dial that set up the existing connection
- The time of a path switch to this connection network path for an existing connection

If this path still has the lowest weight of any available path to the partner node, any attempt to redial the partner node will continue to try the path over this particular VRN, which is likely to result in failure until the underlying problem with the path is corrected.

EE connection network reachability awareness detects the dial failure or connection INOP for the connection over an Enterprise Extender connection network and prevents that specific path to the partner node from being used for a period of time. The path through the Enterprise Extender VRN to the unreachable partner is unidirectional and consists of three components:

- The node on the origin side of the virtual node, which is the node that detects the problem
- The virtual node defined by the Enterprise Extender connection network
- The unreachable partner node on the destination side of the virtual node

It is possible for the path in one direction between two nodes and through the VRN to be usable for route selection, while the path in the other direction is not usable. This function can detect that distinction and when new sessions are established and HPR path switches occur, it allows routing in the direction that is usable, while preventing the path in the direction that is not usable from being selected.

Unreachable time (UNRCHTIM) considerations

With EE connection network reachability awareness, you can use a start option, UNRCHTIM, to indicate that paths to partner nodes over Enterprise Extender VRNs should not be used for route selection for a period of time after initial dial failures or connection INOPs; this provides time for the underlying connection problem to be corrected. An alternate route to the partner node is selected, if an alternate route is available, when new sessions are established and HPR path switches occur. If an alternate route is not available, the session or HPR path switch fails with sense code 80140001. UNRCHTIM can also be coded on a PORT or GROUP definition statement that defines a connection network (with the VNNAME and VNTYPE operands) in an EE XCA major node. This enables a specific connection network to have a different unreachable time value than the value specified on the UNRCHTIM start option. See

the UNRCHTIM start option in [z/OS Communications Server: SNA Resource Definition Reference](#) for more information.

The ramifications of the unreachable time value you choose are important. The results can vary depending on the duration of failures that occur. A value that is appropriate for a short-term failure might not be appropriate for a long-term failure. For example, if the value is set very low and a long-term outage occurs, you introduce the performance impact of added network flows to communicate the unreachable partner information to all nodes each time a path is determined to still be unreachable. In this case, many session or HPR path switch failures can still occur, even though an alternate route is available. Alternatively, if the value is set high and a short-term outage occurs, any paths to unreachable partners detected at that time will remain unavailable for use until the unreachable time expires or until the unreachable partner information is cleared with the MODIFY TOPO,FUNCTION=CLRUNRCH command.

You can change the UNRCHTIM start option value while VTAM is running with the MODIFY VTAMOPTS command. You can also change the unreachable time for a specific Enterprise Extender connection network using the VARY ACT,UPDATE=ALL command with the Enterprise Extender XCA major node after changing UNRCHTIM on the PORT or GROUP definition statement. However, the unreachable time value is not changed for paths to unreachable partners that were determined to be unreachable before the value was modified. When the existing unreachable time value for a path expires, the new unreachable time value will be used if the partner node through that path is again determined to be unreachable.

Rule: The unreachable time for an Enterprise Extender connection network should be set to the same value in all network nodes. If the values are different, the unreachable time used for the Enterprise Extender connection network will be unpredictable. For global connection networks between partner companies, it is practical for unreachable times to be the same but it is not mandatory.

An end node that has detected an unreachable partner path notifies its network node server (NNS) of the unreachability so that the NNS can use that information when calculating routes for the end node. Beginning with z/OS V1R8 Communications Server, the NNS will override the unreachable time value sent by the EN using its unreachable time value for the associated connection network, if configured, or the value specified in its UNRCHTIM VTAM start option. With this enhancement, changing the unreachable time value for an Enterprise Extender connection network will no longer require you to update every attached end node. Only the network nodes and those attached end nodes whose NNS is running z/OS V1R6 or V1R7 Communications Server will need to be updated with the new value. You must still configure an unreachable time value on every end node to activate the EE connection network reachability awareness function.

Performance consideration

By default, the UNRCHTIM start option has a value of 0 for all virtual routing nodes; therefore, EE connection network reachability awareness is disabled. While EE connection network reachability awareness can be enabled for all virtual routing nodes defined at this node (by this VTAM) by setting the UNRCHTIM start option to a value in the range 10 - 65535, the preferred mechanism is to set UNRCHTIM at the GROUP level of the EE XCA major node. This setting provides maximum flexibility in choosing the unreachable time value for each VRN, and in selectively enabling and disabling the function for each VRN.

When EE connection network reachability awareness is enabled for a VRN, APPN Topology and Routing Services begins maintaining unreachability records if failures occur while using that VRN. The presence of unreachability records in the topology database increases the number of CPU cycles that are consumed during APPN route calculation. This CPU cycle increase is magnified as the number of VRNs that contain unreachability records increases and as the average number of unreachability records increases for each of those VRNs.

To mitigate the CPU cost of enabling EE connection network reachability awareness, the following considerations apply:

- Typical EE network designs require that only few connection networks be accessible from a single APPN subnet. For network designs that employ more than four EE connection networks that are accessible from a subnet, EE connection network reachability awareness should be enabled for no more than four of those EE VRNs at any given time.

- While the mechanism employed by EE connection network reachability awareness provides the tracking of connection failures on a per-partner basis, failures to many partners typically indicate a pervasive problem with the underlying transport. You can set a limit on the number of unreachable partner records (partner_limit) associated with a VRN on the UNRCHTIM start option. A network node maintains all unreachability records that are reported from its served end nodes and from all other network nodes in the same network. When the count exceeds this limit of unreachable partner records for a VRN, the VRN is considered unusable for any access until enough existing unreachable partner records expire or are manually cleared with the MODIFY TOPO,FUNCTION=CLRUNRCH command. The count of unreachable partner records must drop below 80% of the unreachable partner limit specified or defaulted on the UNRCHTIM start option for the VRN to again be considered for route selection.

EE virtual node route selected after UNRCHTIM is detected

When an Enterprise Extender VRN is selected by VTAM Topology and Routing Services, a DIAL request is made using the dynamic PU. The session remains in ADIALIP status until the DIAL request completes or until there is a timeout on the connection.

If there is a problem with the physical connection underlying the path through the VRN, the dynamic PU eventually receives an INOP because of the timeout. The following messages are received:

```
IST259I  INOP RECEIVED FOR CNV00074 CODE = 01
IST1411I  INOP GENERATED FOR linename
IST1430I  REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT
.
IST1903I  FAILURE OVER vrnname TO CP partner_cpname
IST2050I  THIS PATH WILL NOT BE SELECTED FOR UNRCHTIM = seconds
          SECONDS
IST314I  END
```

Next, the session fails with sense 08010000. This is normal processing.

When a DISPLAY NET,TOPO,LIST=UNRCHTIM command is entered, the path is now shown as unreachable:

```
IST2057I  UNREACHABLE PARTNER INFORMATION:
IST924I  - - - - -
IST2150I  VIRTUAL NODE vrnname - 1 UNREACHABLE PARTNERS
IST2052I  ORIGIN NODE      PARTNER NODE      UNRCHTIM EXPIRES
IST2055I  origin_cpname    dest_cpname      1800s 07:54:10
IST314I  END
```

However, it appears that TRS is still selecting the unreachable path through the Enterprise Extender VRN because sessions continue to fail with sense code 08010000 following the unreachable partner messages that are received for additional dynamic PUs (for example, if the first INOP is for CNV00074, there can be additional INOPs for CNV00075, CNV00076, and so on).

The reason this is occurring is that these sessions were requested, and TRS selected the route through the EE VRN, before the first PU INOP and the detection of the unreachable partner. It can take a significant amount of time from the DIAL request until the timeout occurs. During this time, a number of sessions can be requested. A DISPLAY SESSIONS,LIST=ALL command will show the session status of the session for which a DIAL request is pending as ADIALIP. The session status of the other sessions waiting for this same path will show as PRAV3. Each of these pending sessions will fail. Further displays of the unreachable partner information will show that the EXPIRES time changes each time an INOP occurs.

This is normal processing for the EE connection network reachability awareness function. Any sessions that were pending before the INOP will fail. However, for any future sessions requested after the unreachable partner information is known, TRS will attempt to route around the EE VRN.

EE connection network reachability awareness in a mixed-release environment

EE connection network reachability awareness can coexist in a mixed-release environment. However, if the PLU of a session resides on an NN with a release earlier than V1R6 (or, without EE connection reachability awareness enabled), or on an EN whose NNS is a release earlier than V1R6 (or a node without

EE connection network reachability enabled), then that session will continue to fail as it did before the introduction of this function.

Consider the following configuration:




Figure 43. EE connection network reachability awareness in a mixed-release environment

Assume the PLU resides on NETA.AA, a pre-V1R6 network node or a node without EE connection network reachability awareness enabled. The SLU resides on NETA.A2, a network node that supports EE connection network reachability awareness and has the function enable for NETA.VRN.

The intermediate node, NETA.A1, is also a network node that has EE connection network reachability awareness enabled for NETA.VRN. Both NETA.A1 and NETA.A2 have the unreachable time value for the EE virtual node, NETA.VRN, set to 180 seconds. There is a problem with the connection that underlies the connection network path through NETA.VRN and it has already been detected in NETA.A1 and NETA.A2.

A DISPLAY NET,TOPO,LIST=UNRCHTIM entered at NETA.A1 would show the following output:

```

IST2057I  UNREACHABLE PARTNER INFORMATION
IST924I  - - - - -
IST2150I  VIRTUAL NODE NETA.VRN 2 UNREACHABLE PARTNERS
IST2052I  ORIGIN NODE  PARTNER NODE  UNRCHTIM  EXPIRES
IST2055I  NETA.A1      NETA.A2      180S      07:54:10
IST2055I  NETA.A2      NETA.A1      180S      07:54:12
IST314I  END
  
```

A request is made to VTAM Topology and Routing Services (TRS) in the PLU network node (or if the PLU is on an end node, the NNS of the EN) to select the route that is to be taken to reach the SLU node. So if the PLU is on NETA.AA, which is a pre-V1R6 node, TRS in NETA.AA has no knowledge of the unusable path through NETA.VRN and, assuming that is the best (lowest weight) path between NETA.A1 and NETA.A2, will select that route.

When NETA.A1 becomes aware of the route to be taken, a DIAL request is made across the bad connection using a dynamic PU (these PUs usually default to a name of CNVxxxx). After a period of time, the DIAL request will timeout and the PU will INOP with the following messages:

```

IST259I  INOP RECEIVED FROM CNVxxxx CODE = 01
IST1411I  INOP GENERATED FROM linename
IST430I  REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT
.
.
.
  
```

```

IST1903I  FAILURE OVER vrnname TO CP partner_cpname
IST2050I  THIS PATH WILL NOT BE SELECTED FOR UNRCHTIM = seconds
          SECONDS
IST314I   END

```

The session will fail in NETA.AA with sense code 08010000. In NETA.A1, where the IST1903I and IST2050I messages were received, the EXPIRES time for the path with NETA.A1 as the origin and NETA.A2 as the partner will be reset to 180 seconds after this failure is detected. The EXPIRES time for the path with NETA.A2 as the origin and NETA.A1 as the partner will not be reset because NETA.A2 did not detect a problem in the reverse direction.

A DISPLAY NET,TOPO,LIST=UNRCHTIM entered at NETA.A1 would now show the following output:

```

IST2057I  UNREACHABLE PARTNER INFORMATION:
IST924I   - - - - -
IST2150I  VIRTUAL NODE NETA.VRN - 2 UNREACHABLE PARTNERS
IST2052I  ORIGIN NODE   PARTNER NODE   UNRCHTIM  EXPIRES
IST2055I  NETA.A1       NETA.A2        180S      07:58:45
IST2055I  NETA.A2       NETA.A1        180S      07:54:12
IST314I   END

```

Sessions routed through the VRN by the pre-V1R6 node (or a node without EE connection network reachability awareness enabled) will continue to fail as long as the connection underlying the EE VRN is not working. However, if the PLU resides on any of the nodes with EE connection network reachability awareness enabled (for example, where that node is selecting the route), sessions will be routed around the bad connection until the unreachable time expires. This is true even when the PLU is on NETA.BA, which is not connected to either of the nodes with EE connection network reachability awareness enabled that detect the failure. This node is informed of the unreachable partner information with Topology Database Updates (TDUs). The pre-V1R6 node, NETA.AA, does not recognize the unreachable partner information about the TDUs, but will pass that information about to NETA.BA, the node with EE connection network reachability awareness enabled.

Displaying unreachable partner information

The DISPLAY TOPO,LIST=UNRCHTIM command displays Enterprise Extender connection network unreachable partner information. Unreachable partner information is maintained only by network nodes and, for this reason, this command is only valid on network nodes. Unreachable partner information identifies paths from one node through an Enterprise Extender virtual node to a second node that are considered unreachable because of a connection network dial failure or connection network link INOP condition. When this occurs, the problem path is remembered for the period of time specified as the unreachable time value for that connection network. During this period, the problem path will not be considered for new sessions or HPR path switches.

To display unreachable partner information, you can specify:

- The name of an origin node on an unreachable partner path, using the ORIG operand
- The name of a virtual node on an unreachable partner path, using the VRN operand
- The name of an unreachable partner (destination) on an unreachable partner path, using the DEST operand

The ORIG, VRN, and DEST operands can be used in any combination to control the scope of the unreachable partner information that is displayed. If you omit all three of these operands, all Enterprise Extender unreachable partner information known by the network node is displayed. Depending on the value of the DSPLYWLD start option, wildcard values can be used for the ORIG, VRN, and DEST operands.

You can use the DISPLAY command to help isolate the location of the problem with the underlying connection. For example, if the display output indicates that a specific partner is unreachable through more than one Enterprise Extender VRN, the problem might be on the partner node or a router on the partner node side of the connection. If the display output indicates that different partners are unreachable through one or more Enterprise Extender virtual nodes from the same origin node, the problem might be on the origin node's side of the connection.

Clearing unreachable partner information

After the underlying problems with the Enterprise Extender connection network paths are corrected, you can clear the unreachable partner information to make these paths available for route selection before the unreachable time value would automatically expire. The `MODIFY TOPO,FUNCTION=CLRNRCH` command can be entered on a network node to clear a set of unreachable partner paths for Enterprise Extender connection networks, making those paths available for route calculation (if they were previously unavailable for route calculation as a result of unreachable partner information).

Use the `SCOPE=LOCAL` operand to clear the unreachable partner paths from only the network node on which the command is entered. Use the `SCOPE=NETWORK` operand to clear the unreachable partner paths from all network nodes in the network. (Remember that topology flows do not occur across network boundaries.)

If a problem is corrected with only one of multiple paths associated with the EE connection network and you clear the unreachable partner information for more than just that one path, the first attempt to establish a new session over a path that is still unreachable (or the first time an HPR path is switched to a path that is still unreachable) has the following results:

- The new session or HPR path switch will fail
- New unreachable path information will be created for the unavailable path

Maintaining unreachable partner information in the topology database

Unreachable partner information about a path through an Enterprise Extender connection network is sent to an end node's network node server when the failure is detected in an end node, or broadcast to all adjacent network nodes when the failure is detected in a network node, using control vectors X'44', X'46', X'47', and X'4C' on topology database updates (TDUs). Therefore, it is possible for the unreachable partner information displayed in one node to have an expiration time that is slightly different from the expiration time for the same unreachable path in another node, because of delays in network traffic.

When a CP-CP session is established between two network nodes, there is an exchange of topology information. The unreachable partner information, including the unreachable time value, is sent in the control vector X'4C'. However, it is possible that a significant amount of time has expired between the time that the failure was detected and the time this topology information is sent. To make sure the time the unreachable partner information expires is consistent, the unreachable time value is modified in the control vector X'4C'. When this occurs, the unreachable value will not be displayed in the node that receives the modified unreachable time value. Instead, five asterisks (*****) will be displayed in the UNRCHTIM column on a `DISPLAY NET,TOPO,LIST=UNRCHTIM` command. However, the expiration times in the node detecting the failure and the node receiving the modified unreachable time value should be consistent.

Unreachable partner information that is detected by an end node and reported to the end node's NNS is not maintained by the end node. If the end node switches to a new NNS, a CP-CP session is established between the end node and its new NNS and the end node's topology information is sent to the NNS. However, because the end node is not maintaining unreachable partner information, any unreachable paths that had been detected by the end node and reported to the old NNS are not included in this topology information sent to the new NNS. Therefore, some failures will be necessary when attempting to establish new sessions, or when HPR path switches occur, until the new NNS learns about the unreachable paths.

Unreachable partner information is not copied to the topology data set when a `MODIFY CHKPT` command is issued with `TYPE=ALL` or `TYPE=TOPO` specified, or when a `HALT` or `HALT,QUICK` command is issued to end VTAM. Therefore, some failures are necessary when attempting to establish new sessions, or when HPR path switches occur, in order to refresh the unreachable partner information after restarting VTAM.

Route selection anomalies

When one path to a partner node over an Enterprise Extender VRN is unreachable, it is possible that the alternative route selected might not be the least weight route from the PLU node to the SLU node. [Figure 44 on page 143](#) and [Figure 45 on page 143](#) are examples of configurations in which this can occur.




Figure 44. Connection network reachability example 1

In the configuration shown in [Figure 44](#) on page 143, the route is calculated in the network node where the PLU resides. The weights of the TGs between the network nodes NN1 and NN2, and between the network nodes NN2 and NN3, are each 240. The weight of the TG between each network node and the VRN is 15, making the paths through the VRN preferred over the direct connections between nodes. Assume that NN3 is an unreachable partner on the path from NN1, through the VRN, to NN3, meaning that this route will not be selected for session setup. The other paths through the VRN are still acceptable for route selection.

With this unreachable path eliminated, the next least weight route goes from NN1, through the VRN, to NN2, through the VRN a second time, to NN3. However, to prevent the calculation of circular routes, VTAM Topology and Routing Services does not allow any node to exist more than one time in a selected route. Therefore, the final route is the higher weight route from NN1, through the VRN to NN2, then from NN2 to NN3.




Figure 45. Connection network reachability example 2

In the configuration shown in [Figure 45 on page 143](#), the route is also calculated in the network node where the PLU resides. The weights of the TGs between the network nodes NN1 and NN2, between the network nodes NN2 and NN3, and between network node NN3 and end node EN1, are each 240. The weight of the TG between each node and the VRN is 15, making the paths through the VRN preferred over the direct connections between nodes. Assume that NN2, NN3, and EN1 are all unreachable partners on the paths from NN1 through the VRN, meaning that these routes will not be selected for session setup. The other paths through the VRN are still considered to be usable for route selection.

With the unreachable paths eliminated, the next least weight route goes from NN1 directly to NN2, then through the VRN to EN1. However, because of the way that routing trees are built, VTAM Topology and Routing Services puts the path from NN1 to the VRN on the tree first. It must remain on the tree because it will be checked later, after the tree is completed, to determine whether the endpoint TG from the VRN to EN1 can be used. This means that, because the VRN is already on the tree, it cannot be added again for the path from NN2 to the VRN (because a node is not allowed to exist more than one time in a routing tree, to prevent circular routes). The same is true for the path from NN3 to the VRN. Therefore, the final route is the higher weight route from NN1 to NN2, from NN2 to NN3, and then from NN3 to EN1.

TCP/IP MTU size for EE

To ensure optimal performance, the TCP/IP maximum transmission unit (MTU) size should be greater than or equal to the RTP network layer packet (NLP) size. The MTU size (both IPv4 and IPv6) might change during the life of the EE connection. The MTU can change for the following reasons:

- 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.
- 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. Some reasons for MTU size changes include:
 - 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 in [z/OS Communications Server: SNA Resource Definition Reference](#) for details), and path MTU changes are discovered in the IP network.

[Table 7 on page 145](#) shows connection conditions and related results. When establishing an RTP connection (for a CP-CP session, to transport ROUTE_SETUP GDS variables, or for an LU-LU session) over an EE connection, RTP pipes learn the MTU size when the pipes are being established (RSETUP flows). RTP then segments data to this size when transmitting outbound data.

<i>Table 7. Connection conditions and results</i>	
If	Then
This node is the origin endpoint of the RTP connection	VTAM sets the maximum packet size equal to the lesser of the MTU size or the VTAM maximum data size
This node is an intermediate node or the destination node of an RTP connection	VTAM sets the maximum packet size equal to the lesser of the MTU size, VTAM maximum data size for the next hop, or the value received on the ROUTE_SETUP GDS variable
This node is one of the endpoints of the RTP connection and a change in the EE connection's MTU size occurs.	When VTAM detects this condition (the EE connection's MTU size changes during the transmission of an NLP) the MTU size is altered. This change is specified in message IST2029I when you issue the DISPLAY EE command. Also, if this change alters the permitted NLP size (NLP size cannot be increased beyond the originally negotiated value for the RTP connection) then this change is specified in message IST1511I which is displayed as the result of the DISPLAY ID=rtp-pu command.

If the MTU size is less than 768 bytes, VTAM sets the maximum packet size to 768 (this is the smallest maximum packet size allowed by VTAM for HPR packets). This limitation can cause TCP/IP to fragment but exists because the RTP layer cannot allow the HPR header to be segmented in the RTP layer.

Running EE in constrained or virtualized environments

To proactively prevent congestion, High-Performance Routing (HPR) flow control (responsive-mode ARB) is sensitive to minor variations in round-trip time or unpredictable response times from the partner. This sensitivity can affect installations in at least two ways:

- Frequent swings in round-trip times lead to ARB rate cuts that prevent the RTP pipe from ramping up to speeds consistent with what the media and network should allow. These rate cuts result in degraded throughput and reduced response time.
- These same swings in round-trip times can cause the RTP endpoint to prematurely enter the path switch state. Although this state does not affect availability, it uses CPU cycles and can result in a significant amount of unnecessary message output.

There are a number of potential causes for these variations. A partner host that is CPU-constrained cannot guarantee consistent and timely responses to ARB flows. Furthermore, the increasing trend toward virtualized environments is making these types of problems more likely to occur.

You can use progressive mode ARB to enable HPR to perform reliably and consistently in platforms using virtualization. Progressive mode ARB is an HPR flow control mechanism that improves the performance of HPR in these environments. Progressive mode ARB is available only to single hop HPR pipes that are over TCP/IP, which includes a single physical hop across a two-hop EE virtual routing node (VRN).

You can use the HPREEARB parameter to define which EE connections will enable progressive mode ARB. Specify the HPREEARB parameter on any of the following statements:

- The GROUP or PU definition statement in the switched major node
- The PU definition statement (DYNTYPE=EE) in the model major node
- The connection network GROUP definition statements in the EE XCA major node

When defined, VTAM negotiates with the HPR partner whether progressive mode ARB will be used on the RTP connection. If both partners require to use progressive mode ARB and the HPR pipe is a single hop

over EE (which includes a single physical hop across a two-hop EE VRN), then progressive mode ARB will be used.

Use the VTAM start option HPRPSDLY to reduce the number of unproductive path switches. With this option, you specify the minimum amount of time, in seconds, that all HPR RTP pipes must delay before entering a path switch state that is caused by an unresponsive partner. During this time, the RTP endpoint periodically tries to contact the partner to avoid switching paths. If you specify the value of 0, RTP pipes initiate path switch processing when a predetermined number of attempts to contact the partner have been unsuccessful. Specify the HPRPSDLY parameter on any of the following statements:

- The PU definition statement in the switched major node
- The PU definition statement (DYNTYPE=EE) in the model major node
- The connection network GROUP definition statements in the EE XCA major node

When you specify the EEDELAY setting for the HPRPSDLY parameter in the major node, VTAM calculates the number of seconds that RTP pipes must delay before entering path switch because of an unresponsive partner. The calculated value allows enough time for the EE keep-alive mechanism to end the EE connection should connectivity be lost to the partner. The benefit is that while EE is determining if there is lost connectivity, the RTP layer is not performing unnecessary path switches. When the EEDELAY value is specified for a path switch delay, and EE is the only path to the RTP partner, the HPRPST value should be specified to a value sufficiently large to allow for the EE connection to be re-established. This allows for the RTP pipe and its associated LU-LU sessions to be successfully recovered.

RTP transmission stall operator awareness and recovery support

If HPR packets are being dropped in the network, an RTP endpoint will request the partner RTP endpoint to retransmit the network layer packet (NLP). There are many reasons why the packet may not be delivered. One common problem is the NLP size which the RTP endpoint is using may be outdated or inaccurate. This situation can lead to an RTP endpoint continuously requesting the same NLP be retransmitted. The sending RTP endpoint honors the request by repeatedly retransmitting the same NLP, but the partner does not receive this NLP. When this occurs, the RTP pipe is considered to be in a transmission stall.

VTAM has support to provide operator awareness of an RTP pipe experiencing a transmission stall. When a specific NLP is retransmitted six times without being acknowledged by the RTP partner (as being received), and at least 10 seconds has elapsed since the NLP was first retransmitted, VTAM will notify the operator with message IST2245I as follows:

```
IST2245I XMIT STALL DETECTED FOR RTP puname TO cpname
```

At this point, VTAM provides RTP transmission stall recovery support for all RTP pipes whose path traverses an Enterprise Extender link. When the RTP endpoint is not directly attached to EE, but the RTP pipe consists of multiple hops with one being a REFIFO hop, VTAM assumes an EE hop is in the path. A REFIFO link is one that indicates that packets flowing across it are not guaranteed to be delivered in order. Recovery support consists one of the following two actions:

- First, if the HPR connection traverses an one hop EE or a two hop EE connection network, the RTP endpoint immediately lowers its NLP size to the HPR-architected minimum size of 768 bytes. All outbound data is then resegmented to this value to resolve the transmission stall.
- Second, VTAM starts a path switch between two RTP endpoints. The path switch allows the EE link to provide the RTP endpoints an updated link size for the EE connection. If the stall persists, VTAM issues message IST2246I every 30 seconds until the stall is alleviated. VTAM also lowers its NLP size to the HPR-architected minimum of 768 bytes if a previous path switch did not alleviate the stalled condition. Message IST2246I follows:

```
IST2246I XMIT STALL CONTINUES FOR RTP puname TO cpname
```


If the recovery actions taken by VTAM result in the packets being delivered to the partner, the transmission stall is considered alleviated at which point VTAM issues message IST2247I as follows:

```
IST2247I XMIT STALL ALLEVIATED FOR RTP puname TO cpname
```

At this point normal data flow resumes. After 20 minutes of operating with the reduced MTU size, this RTP pipe's original NLP size is restored in an attempt to increase RTP throughput. If the network will still not allow transport of the larger packets, the stall detection and alleviation process will repeat.

If at any point, a transmission stall extends beyond the time limit specified by the HPRSTALL VTAM start option, VTAM automatically initiates termination of the HPR pipe and issues the following message:

```
IST2253I HPRSTALL TIME EXCEEDED FOR RTP puname TO cpname
```

Load balancing

If you define multiple links to different VRNs (using different IP addresses), you can choose different TG characteristics on these definitions to force sessions using different APPN COS names to flow over different VRNs (using different local IP addresses). The IP network can also be configured to route traffic for different APPN COS names (IP ToS values) using different physical IP interfaces. For example, #INTER traffic could be sent over a Gigabit Ethernet connection and #BATCH traffic could be sent over a slower OSA link.

Using this configuration, you can monitor APPN traffic flows by displaying the RTPs that flow through a specific VRN. You can accomplish this by specifying the FIRSTCP=vrnname operand, the FIRSTTG=tgnumber operand, or both on the DISPLAY RTPS command. This configuration also enables you to monitor IP traffic flows by tracing flows based on source or destination IP addresses.

If your APPN product does not allow you to define multiple VRNs at the same time, you can define two links to the same VRN (using different IP addresses) and force sessions to be load balanced over those connections. You can still monitor APPN and IP traffic by specifying both the FIRSTCP= and FIRSTTG= operands on the DISPLAY RTPS command in order to see what traffic is flowing over what link.

You might also want to use only a single IP address, but still load balance sessions through parallel VRNs. With this configuration, it is not possible to monitor IP traffic over each of these connections, because there is only a single local IP address being used. But you would still be able to monitor APPN session traffic over the parallel routes (VRNs) by using the FIRSTCP= and FIRSTTG= operands on the DISPLAY RTPS command. Furthermore, because only a single IP address is being used, this type of configuration might simplify the coordination of system definitions.

Transmission group profiles

Six sample TGPs are provided in IBMTGPS for Enterprise Extender.

- Enterprise Extender TGs over WAN (EEXTWAN)
- Enterprise Extender TGs over campus networks (EEXTCAMP)
- Enterprise Extender TGs over Fast Ethernet (FASTENET)
- Enterprise Extender TGs over Gigabit Ethernet (GIGENET)
- Enterprise Extender TGs over 10 Gigabit Ethernet (GIGNET10)
- Enterprise Extender TGs over HiperSockets (HIPERSOC)

As with all APPN links, specify a TGP that represents the actual connectivity being used. The sample TGP FASTENET might be a reasonable choice for EE TGs traversing the Internet. If you do not specify a TGP, the default values for EEXTCAMP are used (a capacity value of 4 MB, a propagation delay value of terrestrial, and a security value of unsecure). Consider the actual capacity and propagation delay with your EE links in order to determine the appropriate TGP or link characteristic values to use.

For more information about these TGPs and details on how to use them, see [z/OS Communications Server: SNA Resource Definition Reference](#).

Dynamic reconfiguration

Enterprise Extender provides substantial flexibility by enabling you to use multiple VIPA addresses or define multiple Enterprise Extender connection networks (local or global). Fully using that flexibility frequently requires coding multiple GROUP statements in the XCA major node for EE, and specifying a number of operands, such as IPADDR, HOSTNAME, VNNAME, TGP, and so on. Given all that flexibility, you might need to modify your definitions to accommodate expanding your existing Enterprise Extender deployment.

Using the VARY ACT,UPDATE command with the XCA major node for EE as the target, you can do the following actions:

- Change operands on a currently inactive GROUP statement
- Add or delete LINE and PU statements under a currently inactive GROUP statement
- Add new GROUP statements and remove unnecessary GROUP statements
- Change virtual node definition values on the PORT statement

When you want to change an Enterprise Extender GROUP statement, use the VARY INACT,ID=*groupname* command to deactivate the GROUP and all subordinate resources under the GROUP. Changing the values of connection network operands on the PORT statement requires deactivation of the GROUP statement identified by the VNGROUP operand.

Dial usability - DWACT, DWINOP, KEEPACT, REDIAL, and REDDELAY

To enable EE connections to automatically activate (without having to issue VARY DIAL commands), put both the EE XCA major node and the switched major node in the VTAM configuration list, and specify DWACT=YES on the EE switched PUs. When VTAM initialization completes, the EE connections will fully activate without further intervention. When the connections are activated, they should be automatically recovered after any type of failure (such as an IP network problem or even a loss of the TCP/IP stack itself).

Certain operands automate the recovery of the EE switched PUs and LINEs when initial activation fails, or when a later connection failure is detected. The operands include:

- KEEPACT = YES | NO

Coding the KEEPACT value as YES on the EE LINE indicates that you want VTAM to immediately try to reactivate the LINE after it fails. This is especially useful in instances where the TCP/IP stack fails or is taken down, thereby making all of the EE LINEs inoperable. With KEEPACT=YES, the lines are reactivated automatically when the TCP/IP stack is recovered.

- DWINOP = YES | NO

Setting Dial When INOP (DWINOP) on the switched PU to YES indicates to VTAM that you want the PU to be redialed after a failure.

- REDIAL = 0 - 254 | FOREVER

If DWINOP=YES is specified on the switched PU, then by default VTAM will make up to 3 redial attempts. The REDIAL parameter can be coded on the PATH statement to change the number of redial attempts to a number in the range 0 - 254, or to FOREVER, meaning that VTAM will redial indefinitely.

- REDDELAY = 1 - 1200

By default, VTAM waits 30 seconds between redial attempts. You can adjust that interval to any value in the range 1 - 1200 seconds by coding the REDDELAY parameter on the PATH statement.

When an EE connection is composed of two z/OS EE endpoints, specifying DWACT=YES or DWINOP=YES on both endpoints can lead to both hosts attempting to initiate or recover an EE connection simultaneously. This can produce unnecessary dial collisions, resulting in delayed activation or recovery, or in some cases failed activation or recovery.

Guideline: When coding DWACT=YES, DWINOP=YES, or both, specify it on only one side of the connection to avoid PU busy conflicts.

Tip: When coding DWINOP=YES, also consider coding the REDIAL and REDDELAY operands on the corresponding switched major node PATH statement. The current defaults for these two parameters are REDIAL=3 and REDDELAY=30. With these defaults values, z/OS Communications Server retries for only approximately 120 seconds after an INOP before giving up attempts to contact the partner EE node.

Customization for EE connection network PUs

Use DYNPU=NO on the GROUP statement when you want dynamic PUs to only be created for inbound and outbound connection network connections. Following is a sample of this type of EE XCA major node:

```
XCAEE    VBUILD TYPE=XCA
PORTEE   PORT  MEDIUM=HPRIP      Enterprise Extender
GPEE1    GROUP ANSWER=ON,
          AUTOGEN=(5,LNLV1,PULV1), Create lines and PUs
          CALL=INOUT,
          DIAL=YES,
          DYNVPFX=C1,           Connection network dynamic PU prefix
          DYNPU=NO,            No dynamic PUs (except conn network)
          ISTATUS=ACTIVE,
          TGP=GIGENET,
          VNNAME=NETA.LVRN1,   Connection network CP name
          VNTYPE=LOCAL         Connection network type
```

With the EE XCA major node shown in the example, the following connections are provided for lines in group GPEE1:

- Inbound and outbound connection network links - the dynamically created PU will be named C1nnnnn
- Inbound and outbound non-connection network links when a matching active switched PU is available

Inbound non-connection network links when a matching switched PU is not available will be rejected with messages similar to the following messages:

```
IST680I CONNECTION REQUEST DENIED - ID = SSCP1A PU GEN NOT SUPPORTED
IST1394I CPNAME = NETA.SSCP1A      STATION ID = 0200FF18A16
IST081I LINE NAME = LNLV1004, LINE GROUP = GPEE1, MAJNOD = XCAEE
IST314I END
```

Use DYNPU=YES on the GROUP statement when you want dynamic PUs to be created for inbound and outbound connection network connections and inbound non-connection network connections. Following is a sample of this type of EE XCA major node:

```
XCAEE    VBUILD TYPE=XCA
PORTEE   PORT  MEDIUM=HPRIP      Enterprise Extender
GPEE1    GROUP ANSWER=ON,
          AUTOGEN=(5,LNLV1,PULV1), Create lines and PUs
          CALL=INOUT,
          DIAL=YES,
          DYNVPFX=P1,           Non-connection network dyn PU prefix
          DYNVPFX=C1,           Connection network dynamic PU prefix
          DYNPU=YES,            Dynamic PUs allowed (all connections)
          ISTATUS=ACTIVE,
          TGP=GIGENET,
          VNNAME=NETA.LVRN1,   Connection network CP name
          VNTYPE=LOCAL         Connection network type
```

With this EE XCA major node, the following connections are provided for lines in group GPEE1:

- Inbound and outbound connection network links - the dynamically created PU will be named C1nnnnn
- Inbound and outbound non-connection network links when a matching active switched PU is available
- Inbound non-connection network links when a matching switched PU is not available - the dynamically created PU will be named P1nnnnn

While the defaults values are often sufficient, you can customize EE connection network PUs in the following ways:

- By default, dynamic connection network PUs take a name in the format CNVxxxxx, where xxxxx is a unique value that is generated and concatenated to the default CNV prefix to create the eight-

character PU name. To specify a different prefix (two characters), use the DYNVNPFX start option or the DYNVNPFX operand on the EE XCA GROUP definition.

- You can create a model PU definition to customize the characteristics of dynamically created PUs. Use the DYNTYPE=VN operand for the model PU definition in the model major node.

Cross-subnet routing with global VRNs

With border node architecture, topology information about nodes, TGs, and their characteristics that determine weights, is not passed across subnet boundaries by way of topology database update (TDU) flows. When a session path is determined to cross subnet boundaries, the optimal route (lowest weight) is calculated in each subnet. The complete end-to-end route consists of contiguous route segments which are individually optimal within their native subnet, but not necessarily optimal from end to end. Global virtual routing nodes (GVRNs) can be used to provide a more direct route between two endpoints in different subnets. Figure 46 on page 150 shows a border node configuration with a GVRN. The weights in the figure are those displayed with DISPLAY NET,TOPO,ID=cpname,APPNCOS=cos or DISPLAY NET,TOPO,ORIG=orig,DEST=dest,APPNCOS=cos commands.

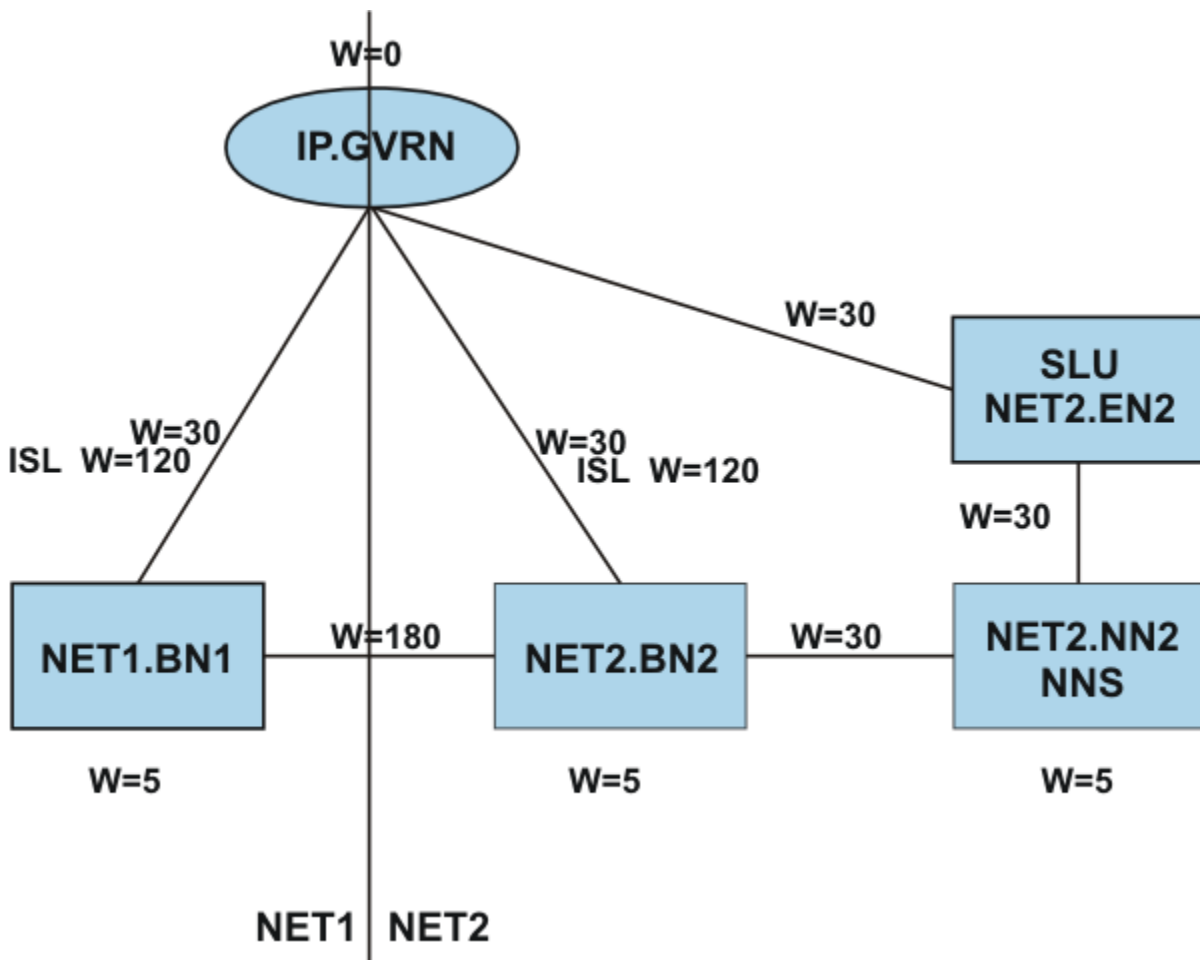


Figure 46. Global VRN with extended border nodes

A TG that originates in a border node and that has a GVRN as the destination can be used across a subnet boundary or used within a single subnet. (An intersubnet link (ISL) between two border nodes cannot be used in this way.) If the TG is used to calculate the session route and the PLU node that calculates the route is in another subnet, then the TG is being used as an intersubnet link (ISL). When the same TG is used within a single subnet, the TG is not considered to be an ISL.

In many ways, the border node in one subnet is treated the same way that an end node is, in that it is considered to be the endpoint of the route calculated in another subnet. For example, in the configuration

shown in Figure 46 on page 150, assume that the PLU is on NET1.BN1 and the SLU is on NET2.EN2. The only information passed to NET1.BN1 on a LOCATE would be:

- ISL TG between the two border nodes
- ISL TG from NET1.BN1 to the global VRN, IP.GVRN (because this TG is being passed cross-subnet, it is considered to be an ISL)
- TG from NET2.EN2 to the global VRN, IP.GVRN.

The ISL between the two border nodes and the ISL TG to IP.GVRN would contain the encapsulated best route from NET2.BN2 to NET2.EN2, but the LOCATE would only include the TG characteristics (in a control vector x'47') for the ISL TGs themselves. Because the weights are determined by the TG characteristics in the x'47' control vectors, NET1.BN1 considers the two TGs to IP.GVRN to have equal weights of 30. Because the two weights are equal, and the weight of the ISL between the two border nodes has a higher weight of 180, one of the two lower weight routes is randomly selected.

For additional information about weight computation see “Example of weight computation” on page 249 and Appendix H, “Forcing an APPN route in a VTAM network,” on page 613.

Because you would normally want your session path to be the two hop route of NET1.BN1->IP.GVRN->NET2.EN2, there is a way to prefer this route over one that transverses an ISL, which might contain additional encapsulated session hops. z/OS Communications Server uses the COSTBYTE TG characteristic to prefer the more direct route. When a TG from a border node to a GVRN is being used as an ISL, the COSTBYTE value of 0 is temporarily changed to 1 (in the x'47' control vector when the TG control vectors are sent to another subnet in a LOCATE) to increase the weight of the TG. If the existing COSTBYTE value assigned to the TG is not 0, no change is made, and the weight will be the same whether the TG is used as an ISL.

The weights used in route calculation are displayed in the output of DISPLAY NET,TOPO,ORIG=orig,DEST=dest,APPNCOS=cos:

```
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP NET2.BN2
IST1357I
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE      CPCP
IST1301I IP.GVRN            21      OPER   INTERM      VALUE  WEIGHT
IST1579I                      -----
IST2241I                      TIME    ISL
IST1163I                      LEFT    WEIGHT
IST1164I                      0      YES    15    120
```

Unlike the ISL to a GVRN, where the COSTBYTE can be temporarily set to 1 when the TG is sent to another subnet in a LOCATE, the COSTBYTE of an ISL between two border nodes is defaulted to 1 when the ISL is activated. The COSTBYTE for this ISL can be changed from the default by specifying a COSTBYTE value on the PU definition for the ISL connection. This can lower the weight of the ISL, causing sessions to be routed through the direct ISL instead of taking the more optimal route through the GVRN.

When you display TG information for an ISL between two border nodes with the DISPLAY NET,TOPO,ORIG=orig,DEST=dest,APPNCOS=cos command, the weight that is displayed in messages IST1300I and IST1301I is the weight of the ISL when it is used cross-subnet:

```
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP NET2.BN2
IST1357I
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE      CPCP
IST1301I NET1.BN1          21      OPER   INTERCLUST  NO    30
IST1579I                      -----
IST2241I                      TIME    ISL
IST1163I                      LEFT    WEIGHT
IST1164I                      0      YES    15    *NA*
```

Guideline: To favor the two hop GVRN route between cross-subnet session partners, you can override the COSTBYTE in a TGP assigned to a direct ISL between two border nodes. This will cause it to have a higher weight than the total weight of a route through a GVRN between the two subnets, considering the ISL weight of a TG between a border node and a GVRN.

TG characteristics (COSTBYTE, COSTTIME, CAPACITY, etc.) can be assigned to a TG either individually or with a TG Profile (TGP), which contains a set of TG characteristics. The TG characteristics of the TGP

assigned to an ISL are specified on the PU statement of the PU that is activated to start the ISL TG. The TG characteristics or the TGP assigned to a TG associated with a GVRN are specified on the PORT or GROUP statement that defines that GVRN in the Enterprise Extender (MEDIUM=HPRIP) XCA major node. See [z/OS Communications Server: SNA Resource Definition Reference](#) for additional information.

Troubleshooting EE problems

Table 8 on page 152 shows some common EE problems you might encounter and suggested solutions to resolve them.

Table 8. Troubleshooting EE problems	
Problem indication	Avoidance method or suggested remedy
<p>Line activation failure:</p> <ul style="list-style-type: none"> • Pending act link • PGAIN failure 	<p>When all of the rules for determining the local static VIPA address (or addresses) for EE are followed, EE line activation could still fail for one of the following reasons:</p> <ul style="list-style-type: none"> • An incorrect TCP/IP stack name was specified with the TCPNAME VTAM start option. • An incorrect source VIPA address was specified with the IPADDR VTAM start option, or on the XCA GROUP definition statement. • An incorrect source VIPA address was resolved from the host name specified as the HOSTNAME VTAM start option, or on the XCA GROUP definition statement. • A host name could not be resolved for the specified source VIPA address. • The IPv4 or IPv6 same host device is not defined or is not started by the appropriate TCP/IP stack <p>If the activation fails, perform one of the following recovery actions:</p> <ul style="list-style-type: none"> • Stop TCP/IP, and activate the correct TCP/IP stack. Then LINE (thus, a PORT) activation should complete normally. • Individually deactivate (force) the LINEs. After the last LINE is deactivated, the PORT deactivates. Thus, the XCA major node is also deactivated. Then, change the TCPNAME, IPADDR, or HOSTNAME start option to the correct value, or update the XCA major node definition to specify the correct IPADDR or HOSTNAME value, and reactivate the XCA major node.

Table 8. Troubleshooting EE problems (continued)

Problem indication	Avoidance method or suggested remedy
<p>Activation failure with the following messages:</p> <pre>IST1411I INOP GENERATED FOR linename IST1430I REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT IST314I END</pre>	<p>This message group is issued when VTAM is not receiving responses to XID requests during activation. It indicates that either the partner is not responding to the request or there are connectivity problems within the IP infrastructure. Some common setup problems that cause this are:</p> <ul style="list-style-type: none"> • IP connectivity has been lost within your network. • EE UDP ports are not defined with consistent values across the network. EE ports should be defined in the range of 12000 - 12004. • EE is not enabled on the remote endpoint. For example, lines and switched PUs are not defined, not activated, or both. • If the EE connection path traverses one or more firewalls, the firewalls must allow UDP traffic to flow for EE ports 12000 - 12004. • If NAT is used in the EE connection path, adhere to these rules: <ul style="list-style-type: none"> – Avoid NAPT. EE does not support NAPT. – When a one-to-one address translation function is performed, the name-to-address resolution mapping for the host name yields the correct NAT address. – If connection network is being used with NAT, you must use HOSTNAME definitions when defining your virtual routing node. <p>Use the D NET, EEDIAG,TEST=YES command. Other helpful commands are PING, TRACERTE.</p>
<p>LU 6.2 sessions do not stay up over EE; your session unexpectedly ends</p>	<p>This problem usually indicates that a limited resource is in use somewhere along the session path.</p> <ul style="list-style-type: none"> • For predefined EE connections, use DISCNT=NO (the default) • For EE VRN-based dynamic connections, consider coding a DYNTYPE=VN model with DISCNT=NO or a delay value of 60+ seconds. <ul style="list-style-type: none"> – Important note for CICS LU6.2 Users: Specifying DISCNT=NO prevents CICS from terminating its sessions at the end of every transaction.
<p>Active EE connection unexpectedly fails with the messages:</p> <pre>IST1411I INOP GENERATED FOR linename IST1430I REASON FOR INOP IS XID OR LDLC COMMAND TIMEOUT IST314I END</pre>	<p>EE connection deactivation because of LDLC timeout. EE periodically tests the EE partner to verify IP connectivity and that the partner is still there. When the tests are unanswered, the EE connection ends with these messages. Some common causes are:</p> <ul style="list-style-type: none"> • The partner unexpectedly ended. • IP connectivity has been lost within your network. • OMPROUTE problems

Table 8. Troubleshooting EE problems (continued)

Problem indication	Avoidance method or suggested remedy
Diagnosis of connection failures or performance considerations	Use the TCP/IP packet trace to analyze Enterprise Extender-related packets flowing to and from a TCP/IP stack on a z/OS Communications Server host. You can use the PKTTRACE statement to copy IP packets as they enter or leave TCP/IP, and then examine the contents of the copied packets. The method of capturing and formatting this trace is described in z/OS Communications Server: IP Diagnosis Guide .
Dial collision problems	See “Dial usability - DWACT, DWINOP, KEEPACT, REDIAL, and REDDELAY” on page 148 for more information about dial collisions.
Poor throughput when using PSRETRY	After each path switch, HPR resets its sending rate to the initial value so frequent path switches can lead to reduced throughput. In particular, setting PSWEIGHT to EQUAL or SAMEROUT can lead to an excessive number of path switches. See z/OS Communications Server: SNA Resource Definition Reference for information about the START option.
Poor HPR throughput over EE with multipath enabled	<p>If MULTIPATH is enabled on the TCP/IP stack, the VTAM start option MULTIPATH is set to TCPVALUE and multiple equal-cost routes exist to the partner EE node, then TCP/IP will round-robin batches of EE packets across each of these routes. If one of these routes cannot reach the partner EE node, then EE may not activate, or if it does, there can be significant performance impacts. See z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures for information about DISPLAY commands.</p> <p>Code or let start option MULTIPATH default to NO. This will disable multipath routing for only EE traffic in the TCPIP stack. This can be updated dynamically using the MODIFY VTAMOPTS command.</p>
High CPU utilization in a branch environment (with lots of EE connections always active)	<p>If you are using many Enterprise Extender connections and these connections are kept indefinitely active (switched PU associated with the EE connection has a DISCNT=NO, either specified or defaulted), the following functions can reduce the CPU utilization by VTAM:</p> <ul style="list-style-type: none"> • HPR alive timer optimization - This function by default is enabled and is controlled by the VTAM start option HPREELIV. For full details of this function, see “HPR ALIVE timer optimization for Enterprise Extender” on page 134. • LDLC Keep-Alive reduction - This function requires you to specify an operand for the LIVTIME Enterprise Extender PORT option. For full details of this function, see “Enterprise Extender LDLC keep-alive reduction” on page 134.

Table 8. Troubleshooting EE problems (continued)	
Problem indication	Avoidance method or suggested remedy
EE connections through the connection network are not rerouting to an alternate path	<p>If the EE connection network path has the lowest weight of any available path to the partner node, any attempt to redial the partner node will continue to try the path over this particular VRN. This is likely to result in failures until the underlying problem with the path is corrected.</p> <p>EE connection network reachability awareness is designed to detect the dial failure or connection INOP for the connection over an Enterprise Extender connection network and prevent that specific path to the partner node from being used for a period of time.</p> <p>Use the EE connection network reachability awareness function to indicate that the path to a partner node over an Enterprise Extender VRN should not be used for route selection for a period of time after the initial dial failure or connection INOP, providing time for the underlying connection problem to be corrected. This function can be enabled in the following manner:</p> <ul style="list-style-type: none"> • Specify the UNRCHTIM start option • Specify the UNRCHTIM operand on either the EE XCA major node PORT or GROUP definition statements
A new EE connection is established between you and a partner company but sessions cannot be established.	<p>This is generally because the firewalls are not allowing UDP traffic on all EE ports. The firewall must allow UDP traffic both inbound and outbound on all five EE ports (12000 - 12004). To assist in diagnosing new EE connection problems, use this command: D NET, EEDIAG, TEST=YES. See z/OS Communications Server: SNA Operation for more information.</p>
<p>A new EE connection fails with the EE health verification failure message:</p> <pre>IST2330I EE HEALTH VERIFICATION FAILED FOR puname AT time</pre>	<p>This is generally because the firewalls or intermediate routers are not allowing UDP traffic on all EE ports. Any firewall or intermediate router must allow UDP traffic both inbound and outbound on all five EE ports, typically 12000 - 12004.</p> <p>To assist in diagnosing new EE health verification failures, use the following command:</p> <pre>D NET, EEDIAG, TEST=YES</pre> <p>See z/OS Communications Server: SNA Operation for more information. Review the Enterprise Extender Connectivity Test output for any unsuccessful test results. See DISPLAY EEDIAG,TEST=YES in z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures for information about analyzing the test output.</p>

Table 8. Troubleshooting EE problems (continued)

Problem indication	Avoidance method or suggested remedy
<p>EE health verification fails on active connections. The following eventual action message is issued on the console:</p> <pre>IST2323E EE HEALTH VERIFICATION FAILED FOR ONE OR MORE CONNECTIONS</pre>	<p>Use the <code>DISPLAY NET,EE,LIST=EEVERIFY</code> command to determine which EE connections are experiencing EE health verification failures. Message IST2325I is displayed for each line or PU, which failed health verification on the most recent LDLC probe to its remote partner. Use the <code>DISPLAY NET,ID=linename</code> or <code>puaname</code> command to get more information that includes local and remote IP addresses. Determine the network connectivity problems between this node and the remote partners. Use the <code>DISPLAY NET,EEDIAG,TEST=YES</code> command to determine the reason of the failure. See DISPLAY EEDIAG,TEST=YES in z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures for more information.</p>
<p>A new EE connection comes up with the EE health verification failure not supported message:</p> <pre>IST2342I EE HEALTH VERIFICATION NOT SUPPORTED BY puname</pre>	<p>During the activation of the EE connection, VTAM sends Logical Data Link Control (LDLC) probes to the remote partner to determine whether all five ports are accessible. VTAM does not receive a response to any of the LDLC probe requests. VTAM continues with the activation of the EE connection between this node and the remote partner. Because VTAM receives no replies to its LDLC probe requests, VTAM determines that the remote partner does not support EE health verification.</p> <p>If EE health verification is required for this PU, contact the remote PU owner about upgrading the PU to support EE health verification probes.</p> <p>If you think that EE health verification supported by the remote partner, use the following command:</p> <pre>D NET,EEDIAG,TEST=YES</pre> <p>See z/OS Communications Server: SNA Operation for more information. Review the Enterprise Extender Connectivity Test output for any unsuccessful test results. See <code>DISPLAY EEDIAG,TEST=YES</code> in z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures for information about analyzing the test output.</p>
<p>The EE connection link terminates because of XID or LDLC timeout.</p>	<p>Consider tuning the LDLC parameters as described in “When does the EE connection go away?” on page 131. Also consider using the <code>D NET, EEDIAG, SRQRETRY</code> command.</p>
<p>The RTP pipe fails to successfully path switch even though an alternate link is available</p>	<p>Because of a problem with the EE connection, an HPR pipe attempts to pathswitch but fails to connect with a message that no alternate routes are available. Ensure that values in the HPRPST start option are all greater than the EE link inoptime. See “When does the EE connection go away?” on page 131 for more information.</p>

Table 8. Troubleshooting EE problems (continued)	
Problem indication	Avoidance method or suggested remedy
Excessive path switch messages (IST1494I) flooding the system console log during large network outage.	Enable the HPR path switch message reduction function with the HPRPSMSG start option. For more information about the HPRPSMSG start option, see z/OS Communications Server: SNA Resource Definition Reference .
You cannot determine the APPNCOS name associated with an RTP <i>puname</i> that unexpectedly deactivates.	Enhance the HPR activation and deactivation messages by setting the HPRITMSG start option to the value ENHANCED. Now, when an RTP is deactivated, you can locate the IST1488I (deactivation) message group on the system console log. You can then find the associated APPNCOS in messages IST1962I, IST1963I, IST1964I, or IST1965I. For more information about the HPRITMSG start option, see z/OS Communications Server: SNA Resource Definition Reference .
<p>RTP transmission stalls repeatedly for pipes which use Enterprise Extender. The following messages are displayed:</p> <pre>IST2245I XMIT STALL DETECTED FOR RTP puname TO cpname</pre> <p>If the stall persists, VTAM issues the following message every 30 seconds:</p> <pre>IST2246I XMIT STALL CONTINUES FOR RTP puname TO cpname</pre> <p>If the transmission stall clears, VTAM issues the following message:</p> <pre>IST2247I XMIT STALL ALLEVIATED FOR RTP puname TO cpname</pre> <p>If the transmission stall extends beyond the time limit specified by the HPRSTALL VTAM start option, VTAM automatically initiates termination of the HPR pipe and issues the following message:</p> <pre>IST2253I HPRSTALL TIME EXCEEDED FOR RTP puname TO cpname</pre>	<p>If path MTU discovery is enabled for IPv4 or IPv6 Enterprise Extender connections, and firewalls are used in the configuration, verify that the firewalls are configured to allow ICMP errors to flow on all hops of the connection.</p> <p>If problems persist, you may consider disabling path MTU discovery for IPv4 Enterprise Extender connections. This can be done by specifying PMTUD=NO in the appropriate ATCSTRxx VTAM start list or on the VTAM START command. Optionally, when VTAM is active, issue the MODIFY <i>procname</i>,VTAMOPTS, PMTUD=NO command.</p> <p>If path MTU discovery is not enabled for IPv4 Enterprise Extender connections, but you still suspect this is an MTU issue, you may consider limiting the maximum packet size which Enterprise Extender will transmit. Consider specifying the MTU operand available on these major nodes:</p> <ul style="list-style-type: none"> • 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.

Chapter 7. OSA-Express

This topic provides an overview of OSA-Express for z/OS Communications Server. An OSA is an integrated IBM Z hardware feature that combines the functions of an IBM Z I/O channel with the functions of a network port to provide direct connectivity between IBM Z applications and their clients on the attached network.

For additional information about planning for and installing OSA-Express, see either the OSA-Express customer's guide and reference in your hardware library or the <https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-z-systems-express-customers-guide-reference>.

OSA-Express overview

OSA-Express is an integrated hardware feature that provides direct connection to clients on local area networks (LANs). The OSA-Express feature plugs into an I/O slot just like a channel card. Fiber-optic cable connects the OSA-Express feature to the network.

OSA-Express supports direct attachment to Ethernet where clients communicate using TCP/IP. For details on using OSA-Express for TCP/IP, see the [z/OS Communications Server: IP Configuration Reference](#). For TCP/IP, OSA-Express can be configured to use IBM Queued Direct I/O (QDIO) architecture to eliminate the need for channel control words (CCWs) and interrupts, resulting in accelerated TCP/IP data packet transmission.

Defining an OSA-Express device to z/OS Communications Server using QDIO

To define an OSA-Express device to z/OS Communications Server using queued direct I/O (QDIO), first you need to define a QDIO transport resource list element (TRLE). This is shown in the following example.

```
TRLHYDRA VBUILD TYPE=TRL
*****
*   TRANSPORT RESOURCE LIST FOR OSA-EXPRESS
*****
*      10      16
HYD1    TRLE  LNCTL=MPC,                *
          READ=(2EC0),                  *
          WRITE=(2EC1),                  *
          MPCLEVEL=QDIO,                 *
          DATAPATH=(2EC2,2EC3),          *
          PORTNAME=(HYD1G1),             *
          PORTNUM=1
```

In this example, MPCLEVEL=QDIO indicates that the direct I/O interface is used. Control data is transmitted across the one READ and one WRITE device. Normal data is transmitted across a DATAPATH device, where each ULP is assigned its own channel unit address. If you define multiple VLANs to the same OSA, you need to configure a separate INTERFACE definition in TCP/IP for each VLAN and a separate DATAPATH device is needed for each of these interfaces. When MPCLEVEL is QDIO, HPDT MPC is also used for this connection across the control channels. The port number (PORTNUM) specifies which physical port on an OSA-Express is to be used for this QDIO device. For OSA-Express and OSA-Express2, only one port, port number zero, is supported for each CHPID. For OSA-Express3 or later, multiple ports are supported on each CHPID. Within a single logical partition (LPAR), each port must have its own unique TRLE definition, with unique read, write, and datapath channel unit addresses. The port name (PORTNAME) is the name that will be assigned to this port. All z/OS LPARs sharing the same port of an OSA-Express must define the same PORTNAME and PORTNUM on their TRLE definitions representing that OSA-Express port.

See [z/OS Communications Server: SNA Resource Definition Reference](#) for more information about coding a TRLE statement.

To define the OSA for IPv4, you must next define an INTERFACE statement in your TCP/IP profile. This is illustrated in the following example.

```
INTERFACE LHYDRA
DEFINE IPAQENET CHPIDTYPE OSD PORTNAME HYD1G1
IPADDR 10.42.1.1/16
VMAC ROUTEALL
READSTORAGE GLOBAL
INBPERF DYNAMIC WORKLOADQ
```

In this example, the PORTNAME value (HYD1G1) must be the port name in the TRLE definition. IPAQENET defines an IPv4 Ethernet interface.

To define the OSA for IPv6, you need an INTERFACE statement in the TCP/IP profile. This is illustrated in the following example.

```
INTERFACE HYD1G16 DEFINE IPAQENET6 PORTNAME HYD1G1
```

In this example, the PORTNAME value (HYD1G1) must be the port name in the TRLE definition. IPAQENET6 defines an IPv6 Ethernet interface.

See [z/OS Communications Server: IP Configuration Reference](#) for more information about coding the INTERFACE statements.

To continue defining an OSA-Express device, perform the following steps:

1. Configure and vary the physical devices online.
2. Activate the TRL deck.
3. Define and start the interface using TCP/IP.

Note: Datapath channels for OSA-Express should not be varied offline while VTAM is active.

When z/OS Communications Server participates in an ensemble environment, QDIO TRLE definitions are dynamically generated for connectivity to the intraensemble data network (IEDN) (CHPID type OSX) if the QDIO interface is defined with CHPIDTYPE OSX and the CHPID parameter. For each dynamic TRLE for OSX, the TRLE name is IUTXT0xx and the corresponding port name is IUTXP0xx (where xx is the configured CHPID parameter). QDIO TRLE definitions are also dynamically generated for connectivity to the intranode management network (CHPID type OSM) if OSM CHPIDs are configured. VTAM searches for two OSM CHPIDs when the TCP/IP stack is initialized, and dynamically generates a TRLE for each CHPID. The TRLE names are IUTMT0xx and the corresponding port names are IUTMP0xx, where xx is the OSM CHPID that is found. See [“Resources automatically activated by VTAM” on page 483](#) for more information about the naming convention of these TRLEs and the number of DATAPATH devices that are dynamically defined.

Tip: If the values that are chosen for the dynamically created OSX TRLE definitions do not suit your needs (for example, you need more than 17 DATAPATH devices), you can define your own QDIO TRLE and configure the QDIO interface with the PORTNAME parameter. See [INTERFACE-IPAQENET OSA-Express QDIO interfaces](#) and [INTERFACE -- IPAQENET6 OSA-Express QDIO interfaces statement](#) in [z/OS Communications Server: IP Configuration Reference](#) for more information.

Restrictions:

- If the QDIO interface that is used for connectivity to the intraensemble data network is defined with the PORTNAME parameter, you must define a QDIO TRLE definition. The PORTNAME parameter value on the TRLE definition must match the PORTNAME parameter value. For information about restrictions on the TRLE name, see [Transport resource list major node](#) in [z/OS Communications Server: SNA Resource Definition Reference](#).
- If z/OS Communications Server runs as a guest on z/VM® and guest LAN definitions is used for connectivity to the intraensemble data network, you must define the QDIO interface with the PORTNAME parameter.

QDIO TRLE definitions are always dynamically generated for connectivity to the intranode management network (CHPID type OSM). For each dynamic TRLE for OSM, the TRLE name is IUTMT0xx and the

corresponding portname is IUTMP0xx (where xx is the value of the OSM CHPID). See [TCP/IP in an ensemble](#) in [z/OS Communications Server: IP Configuration Guide](#) for more information.

OSA routing

For QDIO devices, z/OS Communications Server and OSA-Express provide functions that control how incoming unicast datagrams are routed, especially when the OSA-Express is shared by multiple TCP/IP instances.

When TCP/IP activates a QDIO device, each TCP/IP registers each of its home IP addresses with OSA-Express. (TCP/IP also dynamically registers any updates to its set of home IP addresses with OSA-Express.) This allows OSA-Express to route datagrams destined for a registered IP address to the correct TCP/IP instance.

However, when packets are received for IP addresses that are not registered by any TCP/IP stack, the device needs a method for determining which stack, if any, should receive the packet. Two functions are available to accomplish this; OSA-Express Virtual MAC (VMAC) routing, and primary and secondary routing.

OSA-Express virtual MAC (VMAC) routing

If multiple TCP/IP instances are sharing an OSA-Express, the preferred method of routing is to define or generate a VMAC for each stack for each protocol being used (IPv4 or IPv6). For IPv4, this results in the OSA-Express using the VMAC address rather than the physical "burned in" MAC for all ARPs sent for that TCP/IP stack's registered IP addresses, and using the VMAC as the source MAC address for all packets sent from that stack. In this way, all routers on the same LAN as the OSA-Express will use only the VMAC address as the destination for all packets destined for that specific TCP/IP stack. From a network routing perspective, the OSA-Express with this VMAC will appear as a "dedicated" device to that TCP/IP stack.

This simplifies a shared OSA configuration significantly. The OSA-Express knows by VMAC address exactly which stack should receive a given packet. Even if the IP address is not registered with the OSA-Express, if the packet is destined for that VMAC then the router has determined which stack should be the intermediate router, and the OSA can forward the packet directly to that stack. The capability is given for a stack to indicate to the OSA it only wishes to receive packets to registered IP addresses.

The above simplification is true for IPv6 as well. TCP/IP will use the VMAC address for all neighbor discovery address resolution flows for that stack's IP addresses, and will likewise use the VMAC as the source MAC address for all IPv6 packets sent from that stack. Again, from a network perspective, the OSA-Express with this VMAC will appear as a "dedicated" device to that stack.

The VMAC address may be defined in the stack, or if required may be generated by the OSA. If generated by the OSA, it is guaranteed to be unique from all other physical MAC addresses and from all other VMAC addresses generated by any OSA-Express feature.

The following figure illustrates how OSA-Express virtual MAC routing works:

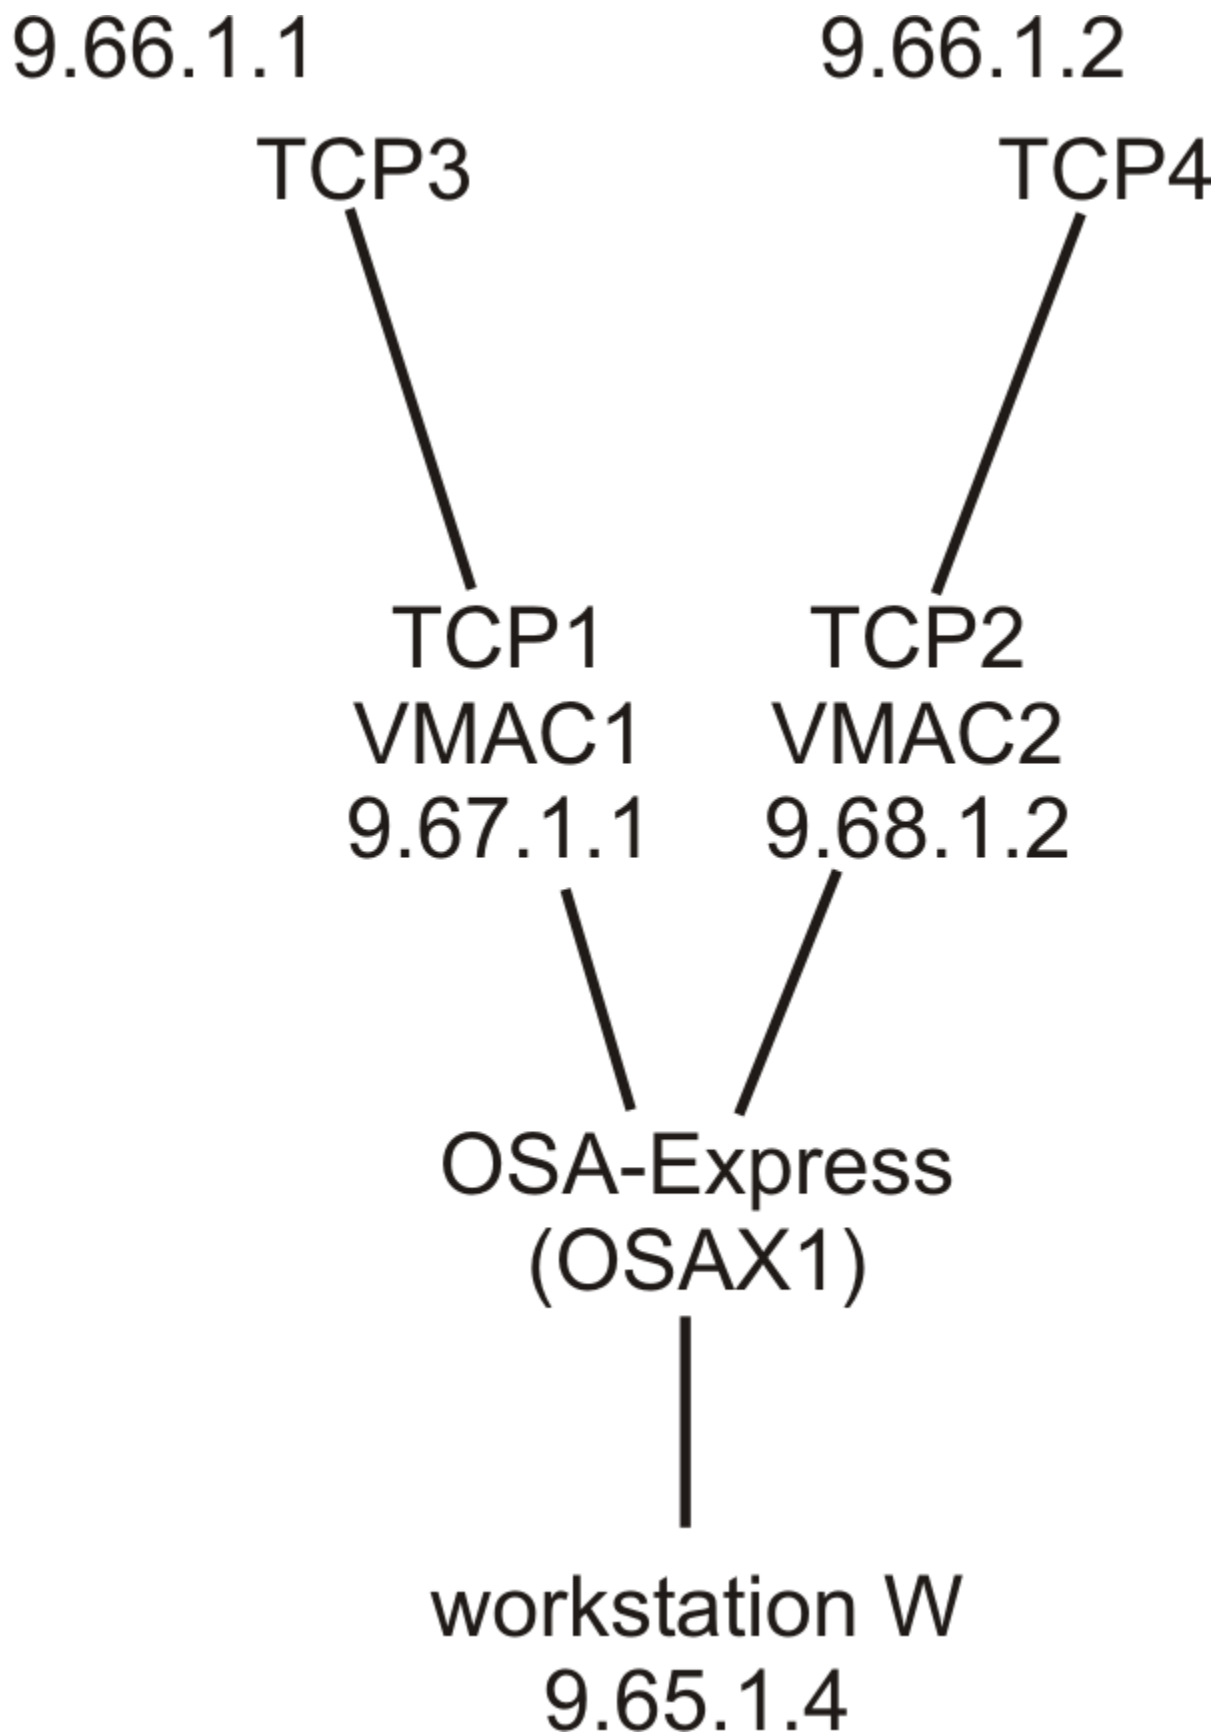


Figure 47. OSA-Express virtual MAC routing

In this example, both TCP1 and TCP2 activate OSA-Express device OSAX1. TCP1 is defined with VMAC1 and TCP2 is defined with VMAC2. TCP1 registers each of its home IP addresses (just 9.67.1.1 in this example) to the OSA-Express device when the device is started on TCP1. The OSA-Express device indicates in all ARP processing that IP address 9.67.1.1 can be reached by using VMAC1. Similarly, TCP2 registers each of its home IP addresses (just 9.68.1.2 in this example) to the OSA-Express device when the device is started on TCP2, and again the OSA-Express device indicates in all ARP processing that IP address 9.68.1.2 can be reached by using VMAC2.

If device OSAX1 receives any datagram destined to VMAC1, the OSA-Express routes that datagram to TCP1. So if routers determined either that IP address 9.66.1.1 or 9.67.1.1 could be reached by VMAC1, TCP1 receives packets destined to those addresses. Likewise, if device OSAX1 receives any datagram destined to VMAC2, the OSA-Express device routes those datagrams to TCP2. Therefore, if routers determined that either IP address 9.66.1.2 or 9.68.1.2 could be reached by VMAC2, TCP2 receives packets destined to those addresses.

Rule: If VMACs are defined in the stack, they should be defined as locally administered MAC addresses, and should be a unique address for the local LAN on which they reside. For IPv4, the VMAC is defined on the IPAQENET type INTERFACE statement representing the shared OSA-Express. For IPv6, the VMAC is defined on the IPAQENET6 type INTERFACE statement representing the shared OSA-Express.

Guidelines:

- If the OSA is configured for both IPv4 and IPv6 for a stack, then the stack must use one VMAC on the INTERFACE statement for IPv4 usage, and one VMAC on the INTERFACE statement for IPv6 usage.
- A VLAN id may be associated with an OSA-Express link or interface defined with a VMAC.

You can define multiple interfaces to the same OSA for the same IP version by configuring multiple VLANs. To do this, configure a unique VLAN ID for each IPv4 interface over the OSA-Express and a unique VLAN ID for each IPv6 interface over the OSA-Express and configure (or generate) a unique VMAC for each INTERFACE statement.

See [z/OS Communications Server: IP Configuration Reference](#) for how to define VMACs and VLANs in TCP/IP. See [z/OS Communications Server: IP Configuration Guide](#) for more considerations when using an OSA-Express configured in QDIO mode and for information about using multiple VLANs.

See <https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-z-systems-express-customers-guide-reference> for the level of OSA-Express2 that supports OSA-Express Virtual MACs.

Primary and secondary routing

The following information applies to IPv4 datagrams. For primary and secondary router specifications for IPv6 datagrams, see the INTERFACE statement contained in the [z/OS Communications Server: IP Configuration Reference](#).

If only one TCP/IP instance is using the OSA-Express, or when multiple TCP/IP instances are using the same OSA-Express device, and you wish all instances to share the same physical MAC address of the device, you can optionally have the OSA route packets to unregistered IP addresses by designating a TCP/IP instance as the primary router or secondary router. For a given OSA-Express, only one TCP/IP instance can be registered as the primary router and one or multiple TCP/IP instances can be registered as the secondary router, depending on the level of the OSA-Express. (See the description of the PRIROUTER and SECROUTER parameters on the [INTERFACE-IPAQENET OSA-Express QDIO interfaces](#) in [z/OS Communications Server: IP Configuration Reference](#) for details on how to designate a primary or secondary router.) If you define a TCP/IP instance as a primary or secondary router, you must also enable IP forwarding (using IPCONFIG DATAGRAMFWD in the TCP/IP profile).

When the device receives a datagram, OSA-Express routes the datagram as follows:

- If the datagram is destined for a registered IP address, OSA-Express routes the datagram to the TCP/IP instance that registered the IP address.
- If the datagram is destined for an unregistered IP address, OSA-Express routes the datagram to the TCP/IP instance that is registered as the primary router. If no TCP/IP instance is registered as the primary router, OSA-Express routes the datagram to one of the TCP/IP instances that is registered as

the secondary router. If no TCP/IP instance is registered as a secondary router, OSA-Express discards the datagram.

Note: If a device becomes inactive for any reason, then OSA-Express unregisters any IP addresses that had been registered.

With QDIO, each TCP/IP instance using the device automatically registers its entire home list (including VIPAs) with OSA-Express, and OSA-Express automatically creates and initializes the OSA Address Table (OAT). There is no need for the customer to configure any home IP addresses in the OAT.

The QDIO primary and secondary routing function can provide fault tolerance without using virtual IP addresses (VIPA).

The following example [Figure 48 on page 164](#) illustrates how the QDIO primary and secondary routing works:

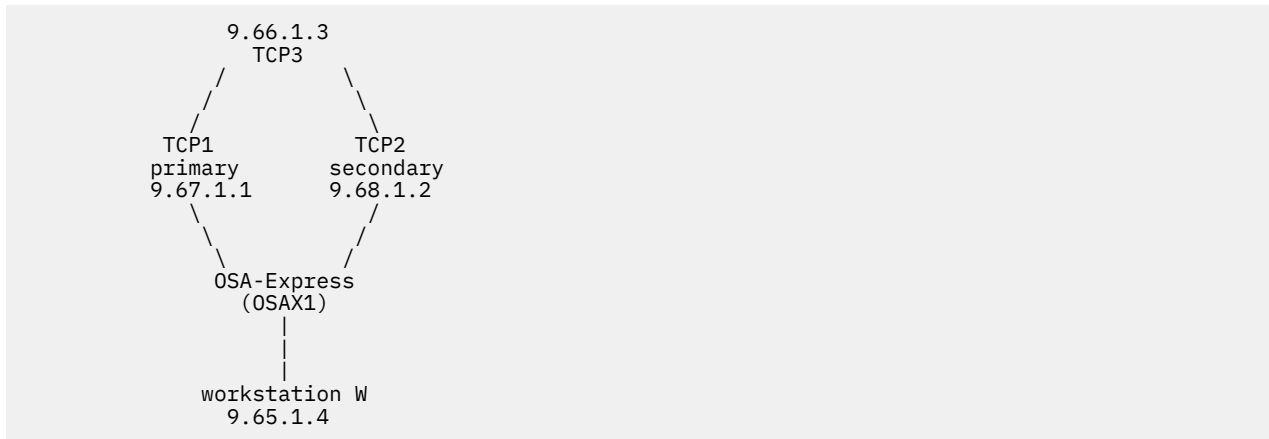


Figure 48. QDIO primary and secondary routing

In this example, both TCP1 and TCP2 activate OSA-Express device OSAX1. TCP1 is defined as the primary router and TCP2 is defined as a secondary router. TCP1 registers each of its home IP addresses (just 9.67.1.1 in this example) to OSA-Express when the device is started on TCP1. Similarly, TCP2 registers each of its home IP addresses (just 9.68.1.2 in this example) to OSA-Express when the device is started on TCP2.

If device OSAX1 receives an inbound datagram destined for 9.67.1.1, then OSA-Express will route the datagram to TCP1. If the device receives a datagram destined for 9.68.1.2, then OSA-Express will route the datagram to TCP2. If the device receives a datagram destined for any other (unregistered) IP address such as 9.66.1.3, then OSA-Express will route the datagram to TCP1 because TCP1 is the primary router. If TCP1 goes down or device OSAX1 becomes inactive on TCP1, then OSA-Express will route such datagrams to TCP2. If the device subsequently becomes active again on TCP1, then OSA-Express will once again route datagrams for unregistered IP addresses to TCP1.

IP traffic from workstation W destined for 9.66.1.3 will go to TCP3 through TCP1 (because the default router is the primary router TCP1). If device OSAX1 becomes inactive on TCP1 for any reason (or TCP1 goes down), then TCP2 (secondary router) becomes the default router and the IP traffic will go to TCP3 through TCP2. If device OSAX1 becomes active again on TCP1, then TCP1 once again becomes the default router and the IP traffic will now go through TCP1.

Note: If you are using static routing, then in order for the above to work, you need to ensure that you have static routes defined in both directions between TCP3 and the workstation (9.65.1.4) through both TCP1 and TCP2.

For more information about primary and secondary routing, see <https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-z-systems-express-customers-guide-reference>.

You can also use VIPA to provide fault tolerance. See information about Virtual IP Addressing (VIPA) in the [z/OS Communications Server: IP Configuration Guide](#) for more details on VIPA.

Outbound priorities

For QDIO devices, z/OS Communications Server and OSA-Express provide a function that assigns a priority value to each outbound datagram and that attempts to provide preferential service to the higher-priority data.

z/OS Communications Server supports four priority values in the range 1–4 for outbound QDIO traffic (with 1 being the highest priority).

TCP/IP uses the first three bits of the Type of Service (ToS) byte in the IP header to determine the outbound priority value for a given datagram. The default mapping of ToS values to priorities is:

ToS	Priority
000	4
001	4
010	3
011	2
100	1
101	1
110	1
111	1

You can use the z/OS UNIX Service Policy Agent to override the default mapping of ToS values to priorities. (See the description of the `SetSubnetPrioTosMask` statement in [z/OS Communications Server: IP Configuration Guide](#) for details on how to override the default priority for a given ToS setting.)

Note: You cannot use a virtual IP address (VIP) as the SubnetAddr value on the `SetSubnetPrioTosMask` statement.

MTU

The maximum transmission unit (MTU) for QDIO gigabit Ethernet devices is 8992. The MTU for QDIO fast Ethernet devices is 1492.

Chapter 8. Defining resources dynamically

You can reduce the number of definitions that must be coded by using the dynamic definition capabilities that VTAM provides. VTAM enables the dynamic definition of switched resources and channel-attached devices. It also enables the dynamic reconfiguration of existing resources and the dynamic change of operands on existing definitions.

The following lists shows where to find information about dynamically defining resources.

- [“Defining switched resources dynamically” on page 167](#)
- [“Dynamic configuration of channel-attached devices” on page 170](#)
- [“Dynamic reconfiguration and change of operands” on page 174](#)

Defining switched resources dynamically

You can dynamically define switched peripheral nodes (including LAN-attached nodes) using:

- Dynamic PU definition (DYNPU operand)
- The dynamic switched definition facility

The following sections describe this process.

Dynamic PU definition (DYNPU operand)

The DYNPU operand enables PUs for type 2.1 peripheral nodes to be created dynamically on the VTAM host accepting an incoming call. Code DYNPU=YES on the GROUP definition statement in an NCP, external communication adapter (XCA), or channel-attachment (CA) major node.

A switched PU is dynamically created when a calling PU cannot be identified by VTAM during a switched dial-in operation. VTAM uses the model PU definition in the model major node to define the characteristics of the PU.

Note: The DYNPU operand is *ignored* if a configuration services XID exit routine is active when a dynamic PU is needed. The exit routine determines whether a PU will be dynamically defined.

Dynamic switched definitions

Using the dynamic switched definition facility, you can dynamically create definitions for types 1, 2, or 2.1 switched dial-in devices, and for leased devices that appear to VTAM to be switched (for example, devices connected by NTRI). You cannot dynamically define type 4 or type 5 devices. To create the dynamic definitions, this facility uses information provided by an installation-supplied exit routine and coded in a model major node.

The dynamic switched definition facility requires model definition statements and a configuration services XID exit routine. VTAM uses the model definitions to create the dynamic switched major node, from which definitions are taken to build dynamic switched devices.

A configuration services XID exit routine is invoked when an unknown device dials in to VTAM. The exit uses IDBLK, IDNUM and other parameters from the XID vector to gather resource definition information. The exit returns to VTAM the model statements, an optional dynamic switched major node name, and the names for the new PU and LUs.

If necessary, VTAM builds a dynamic switched major node. The new devices are added to the major node. If the exit does not specify a major node name, VTAM uses ISTDSWMN. [Figure 49 on page 168](#) illustrates the process of creating resources in a dynamically switched node.

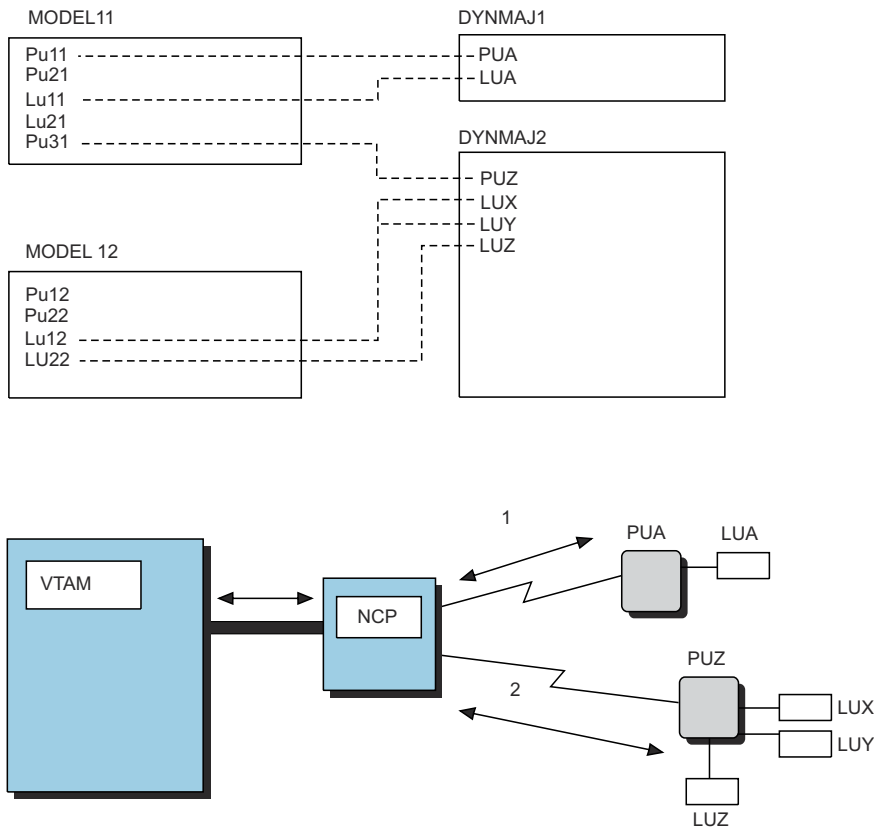


Figure 49. Creating resources in dynamic switched major node

In Figure 49 on page 168, when device PUA dials in to VTAM, PU and LU definitions from MODEL 11 are used by a configuration services exit routine to create resources in an existing dynamic switched major node, DYNMAJ1. When device PUZ dials in to VTAM, PU and LU definitions in MODEL 11 and MODEL 12 are used by the exit routine to create resources in dynamic switched major node, DYNMAJ2. Because dynamic switched major node DYNMAJ2 did not already exist, it, along with its resources, was created by this process.

Configuration services exit routine

A sample configuration services XID exit routine (ISTEXCCS ASSEMBLE) is provided in the sample library (SYS1.SAMPLIB). You can write your own exit routine or use the sample. To code the exit routine to build dynamic resources, follow these steps:

1. Code the exit to supply names for:

- The switched PU and LU resources
- The model definition to be used for the PU or LU
- The dynamic switched major node where the dynamic PUs and LUs are to be created

Note: If the exit does not return a major node name, the default major node ISTDSWMN is used. If the exit returns an unknown major node name, a new dynamic switched major node will be created.

2. Install the exit routine on VTAM to support dynamic switched devices.

3. Activate the routine using the MODIFY EXIT command; otherwise, no dynamic definitions are built, and the device request for connection is rejected.

For information about coding this exit routine, see [z/OS Communications Server: SNA Customization](#).

Model major nodes and model statements

For additional dynamic PU definitions that are possible using the configuration services XID exit, see [Defining a PU and an LU for the configuration services XID exit routine in z/OS Communications Server: SNA Resource Definition Samples](#).

You can code one or more model major nodes, or copy and modify the model major node provided on the sample library (ISTMODEL on SYS1.SAMPLIB). Model major nodes contain the parameters that VTAM uses to create dynamic switched resources.

Define a model major node with a VBUILD definition statement (TYPE=MODEL). The PU and LU definition statements in the model major node define the characteristics of switched peripheral nodes.

Model PU and LU statements are defined for the minor nodes. When the configuration services XID exit routine returns the names of the model resource and the previously undefined resource, VTAM uses these model statements to create the dynamic PUs and LUs. The exit can also identify the name of a dynamic switched major node. There is no relationship between these PU and LU definition statements in a model major node. The exit routine can provide the name of any model PU or LU definition that matches the device characteristics. Therefore, there is no sifting done from the PU definition to the LU definition. The GROUP definition is available only for sifting down to the LU definitions.

Sample model major node

In the following sample you do not need to code LU statements following a PU statement.

MODEL	VBUILD	TYPE=MODEL	
SPUMOD02	PU	ADDR=02, MAXDATA=256, PUTYPE=2, MAXOUT=1, IRETRY=NO, DISCNT=YES, ANS=CONTINUE	MODEL MAJOR NODE PU STATION ADDRESS MAX BYTES PU RECEIVES IN ONE PIU PU TYPE MAX PIUS WITH NO REPLY NCP RETRY DISCONNECT AT SESSION END CONTINUE DURING NCP ANS PU STATION ADDRESS
SPUMOD01	PU	ADDR=01, MAXDATA=256, PUTYPE=2, MAXOUT=1, IRETRY=NO, DISCNT=YES	
GROUP1	GROUP	PACING=(1,1), VPACING=2	PACING LU AND BOUNDARY NODE PACING VTAM AND BOUNDARY NODE
SPUMOD	PU	ADDR=03, MAXDATA=256, PUTYPE=2, MAXOUT=1, IRETRY=NO, DISCNT=YES	
SLUMOD01	LU	LOCADDR=1, MODETAB=MODETAB2	LU LOCAL ADDRESS LOGON MODE TABLE NAME
SLUMOD03	LU	LOCADDR=3, MODETAB=MODETAB2	LU LOCAL ADDRESS
SLUMOD02	LU	LOCADDR=2	LU LOCAL ADDRESS

Note: The PACING and VPACING values coded on the GROUP statement do not sift down to SPUMOD. SPUMOD gets the default values for PACING and VPACING. PACING and VPACING sift down to the LUs from the GROUP, but they would not have sifted down from the PU if they had been specified there.

A PU on a switched line is identified by the device CPNAME or IDBLK and IDNUM. As soon as a dial-in line has been activated and placed in answer mode (during VTAM startup or by the VTAM operator), the physical unit can dial the line number and establish communication with VTAM.

When the dial-in connection is broken, the devices are deactivated and the PU and LU definitions in the dynamic switched major node are deleted. When the last PU within the dynamic switched major node is deleted, the switched major node is also deleted.

Coding DISCNT=YES causes the connection to be broken after all sessions using the line have ended. If you code DISCNT=YES on your model PU statements, consider the cost of the line and the number of session starts and ends this particular device will have.

Restriction: Dynamic definition of switched connections is for dial-in use only. VTAM has no knowledge of the path definitions. Therefore, should the line drop, even if the device is active, VTAM cannot use the definition for dial-out operations. The sole exception to this rule is for a model PU defined for Enterprise Extender with DYNTYPE=EE and DWINOP=YES coded. In this case, if the EE connection is dropped, VTAM will attempt to dial back out to the partner using the remote IP address and remote SAP sent inbound when the connection was established. The configuration restart files are not updated with dynamic definitions.

In environments with an owning and backup host, the backup host can also implement a configuration services XID exit, and dynamically create resource definitions during takeover of NCPs that have dial-in resources connected to them. The backup host exit is provided the same information during takeover that is provided to the exit during a dial-in operation.

Dynamic configuration of channel-attached devices

Dynamic configuration of channel-attached devices (dynamic I/O) enables you to modify your I/O configuration without disrupting your system. You can add, delete, or modify I/O components, such as devices and paths, without having to re-IML or re-IPL.

Use the following items to dynamically configure your channel-attached devices and dynamically create VTAMLST members for new devices:

- Hardware configuration definition (HCD)
- NetView Release 3 and subsequent releases, or a product with a similar interface and services

Note: NetView and TSO/E are necessary for dynamically creating VTAMLST members using the VTAM-supplied sample command lists. To modify the I/O configuration only, you need only MVS and HCD.

Figure 50 on page 171 shows how a channel-attached device is added using dynamic I/O configuration support.

1. Physically attach a device.
2. Add the device to your configuration using hardware configuration definition (HCD), and activate the new configuration using the MVS ACTIVATE command.
3. MVS issues message IOS502I, which causes NetView to start the IOS502I command list.
4. The IOS502I command list calls the ISTDEFIN command list for each device added to the system.
5. The ISTDEFIN command list obtains a description of the new device using the ISTDINFO function module, which obtains the new device information and returns it to ISTDEFIN.
6. The ISTDEFIN command list builds a VTAMLST member for the device. Any information the device cannot provide can be found with IBM- or user-supplied defaults.
7. The ISTDEFIN command list issues a VARY ONLINE command for the device if it is not already online.
8. The ISTDEFIN command list issues a VARY ACT command to activate the device.

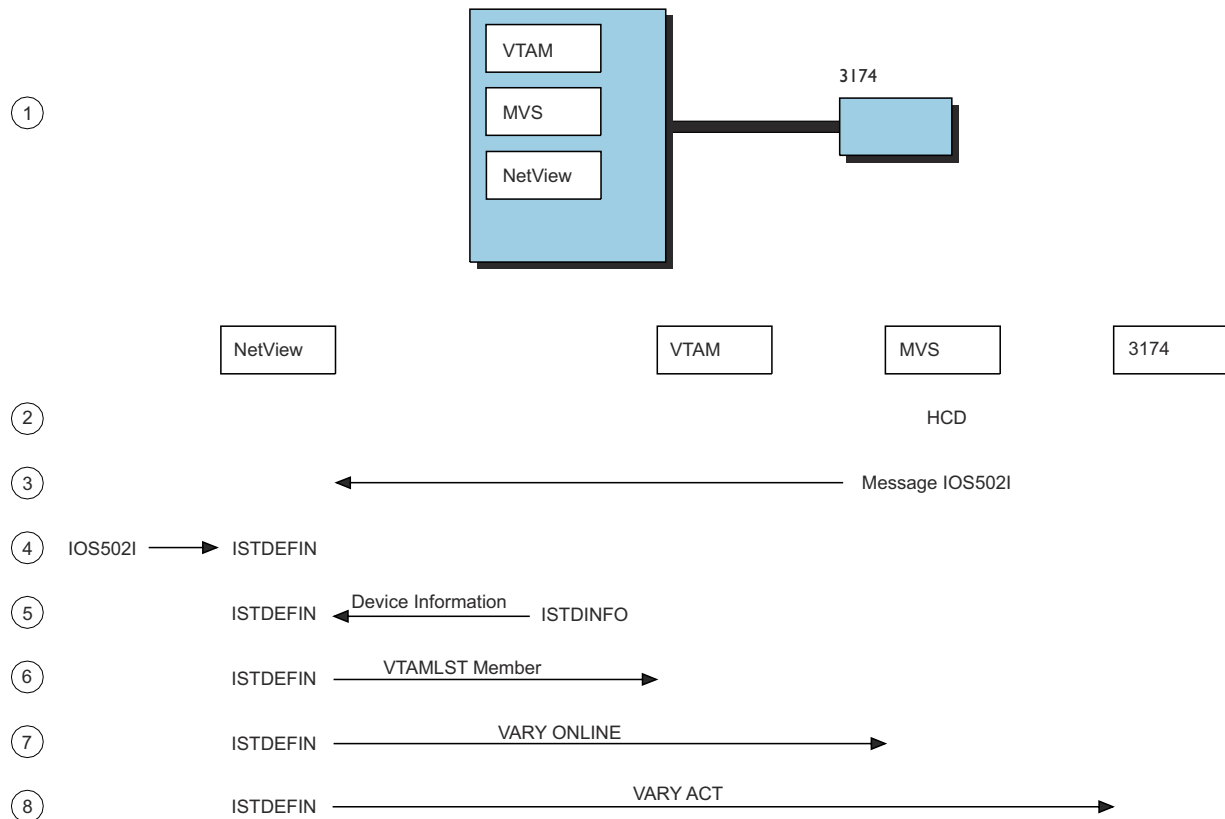


Figure 50. Dynamic configuration of channel-attached device

Installation and preparation

You can use the NetView program or another product that provides similar interface and services to dynamically configure your channel-attached devices.

Procedure

If you are using NetView, take the following steps to prepare it for dynamic I/O configuration:

1. Modify the message automation table so that when MVS issues message IOS502I, NetView can start the IOS502I command list.

The following sample addition to the message automation table logs messages issued by the ISTDEFIN command list, starts the IOS502I command list, routes the message to one operator and logs it:

```
IF MSGID="ISTDEFIN" THEN
  DISPLAY (N)
  NETLOG(Y);
IF MSGID="IOS502I" THEN
  EXEC (CMD("IOS502I") ROUTE (ONE OPER1))
  DISPLAY (N)
  NETLOG(Y);
```

2. Modify the NetView start procedure to point to the sample library (VTAM.SAMPLIB) where the IOS502I and ISTDEFIN command lists reside or move the command lists from the sample library to the library where your other command lists are stored.
3. Activate the message automation table that you changed using the NetView command AUTOMSG MEMBER=tablename.

Results

See the [z/OS Communications Server: New Function Summary](#) for information about data sets needed for dynamic I/O configuration.

Defining your configuration

After preparing NetView or a similar package, use the hardware configuration definition (HCD) to interactively define your I/O configuration changes. (Use of the HCD is described more fully in the MVS library.)

To attach a device to the system:

1. Physically attach the device.
2. Define the device to MVS using HCD.
3. Activate the new configuration using the MVS command ACTIVATE IODF.

MVS then issues message IOS502I, which causes NetView to start the IOS502I command list. The IOS502I command list calls the ISTDEFIN command list for each device added to the system.

Note: Message IOS502I does not display on the MVS console nor on the NetView operator console unless requested using the NetView ASSIGN command.

Building resource definitions

The ISTDEFIN command list obtains descriptions about the new devices using the ISTDINFO function module. The device-specific information that is returned by the ISTDINFO function module varies according to the device.

The ISTDEFIN command list builds channel attachment or local SNA major nodes (PU type 2, 2.1, and 4) for the devices.

The ISTDEFIN command uses IBM-supplied defaults to build the definition statements. [Table 9 on page 172](#) shows the devices and the definitions that the ISTDEFIN command builds.

Table 9. Definitions for dynamically configured devices	
Device	Definition
3705	Channel-attachment major node
3720, 3725	Channel-attachment major node
3745 (type 6 adapter)	Channel-attachment major node
3745 (type 7 adapter with NCP loaded)	Channel-attachment or Local SNA major node
3274	Local SNA major node
3174	Local SNA major node

For a 3745 device that is self-describing (3745 type 7 adapter), the ISTDEFIN command builds either a channel-attachment major node or a local SNA major node to define the device. The definition built is determined by the NCP generation already residing in the 3745. For example, an NCP level that does not support a type 7 adapter causes the ISTDEFIN command to build a channel-attachment major node. Otherwise, the major node built depends on whether the channel is defined as a type 2 (Local SNA major node) or type 4 (channel-attachment major node) in the NCP.

Information needed for defining the major nodes that is not provided by the ISTDINFO function module is defined in templates in the ISTDEFIN command list.

Notes:

1. If the device is self-describing, an NCP being defined must already be loaded.

2. The VTAM using dynamic I/O configuration can only be a data host to an NCP. VTAM does not alter existing PATH definition statements to create a route to the NCP. Do this before you try to dynamically add any devices. Otherwise, the activation of the channel attachment will fail.
3. A local SNA major node can be dynamically created to represent a type 2.1 channel-attached to an NCP. However, no cross-domain resources (CDRSC) definitions for the independent LUs defined under the type 2.1 channel are dynamically created. Also, default adjacent link station lists belonging to the independent LUs are not updated with the dynamically-created PU name.

Using the default naming convention

The IBM-supplied command lists generate names for resource definitions using either an IBM-supplied or user-written naming convention. Names are generated in the following manner:

- The first character is N, identifying the name as being generated by VTAM for dynamic I/O configuration.
- The second character is the subarea number expressed as a single character (0–9 or A–Z). For example, a communication controller in subarea 6 has a second character of 6. A communication controller in subarea 35 has a second character of Z. A communication controller in subarea 36 has a second character of 0, because subarea numbering begins with 1, not 0. The maximum number of subareas that can be defined using the IBM-supplied naming convention is 36.
- The next four characters are the 4-digit device address, represented in hexadecimal.
- In the definition of a single device, a number of resources must be named. The last characters depend on the resource being named:

L

LINE

G

GROUP

D

Data-set name

xx

2-digit hexadecimal local address for the logical unit attached to the cluster controller (3x74)

For example, the name generated for the definition of a device at X'123' in subarea 6 is named N60123D. The LINE name for the device is N60123L.

Customizing the command lists

The IOS502I and ISTDEFIN command lists can run without any modification. However, you can modify the command lists to match the needs of your environment. The command lists are coded in REXX, and the subroutines and functions are documented in the code. For a description of the methods for calling the command lists and the information returned, see [Appendix F, “Command lists: Dynamic configuration of channel-attached devices,”](#) on page 601.

Modifying the resource definition defaults

Information needed for defining the major nodes that is not provided by the ISTDINFO function module is defined in templates in the ISTDEFIN command list. The templates contain information about logon modes, interpret tables, resource names, the number of logical units (SNA cluster controllers only), and other information.

Some information, such as the number of buffers VTAM should allocate for receiving data from the NCP (provided on the MAXBFRU operand of the HOST definition statement), is not always available. You can modify the defaults in the templates to override the IBM-supplied defaults. Edit the ISTDEFIN command list, which shows you the templates with which you can modify the values in the REXX statements.

Note: When logical units are being defined, if the physical unit presents a *logical_terminals* string, only the first LU statement in the template is used.

Modifying the naming convention

You can change the naming convention by modifying the ISTDEFIN command list. A set of directions is provided in the ISTDEFIN command list for supplying a *seed* value to control the values for the first two characters of the device name. For example, if you want to have your device names begin with MY, change the following in the ISTDEFIN command list:

```
seed='MY'                /* user seed value */
```

Thus, the name generated for the definition of a device at X'123' is named MY0123D.

If you want to use the IBM defaults, do not change the *seed* value.

ISTDEFIN replaces the major node if the same name already exists in USER1.AUTO.VTAMLST, so the definition contains the latest device information.

Supplying your own naming convention

If you want to use your own naming convention, change the logic of the ISTDEFIN command list. The functions and subroutines that obtain the system parameters used to build the definitions (such as device number and subarea number) are provided and documented in the command list. You can use these functions to modify the ISTDEFIN command list to use your own naming convention.

Writing your own command lists

The ISTDEFIN and IOS502I command lists are designed to work together. However, you can call these command lists with your own functions and routines.

For example, you can write a command list that is invoked when NetView or VTAM startup completes. This command list can determine all VTAM-supported devices not currently used by VTAM. This command list can check all possible devices and then call the ISTDEFIN command list to build the definitions and activate the devices.

Dynamic reconfiguration and change of operands

Dynamic reconfiguration is the process of adding, deleting, or moving resources within your network configuration without deactivating the affected major node. You can use dynamic reconfiguration for local peripheral nodes, Enterprise Extender XCA major nodes, application major nodes, and CDRSC major nodes.

Dynamic change of operands enables you to modify certain operand values on resource definition statements without deactivating the major node.

Tip: LUs and PUs can be defined dynamically, in which case no system definition is required.

Dynamic reconfiguration and dynamic change of operands are supported for the following resources:

- PUs, LUs, and PATHs can be added and deleted in a switched major node
- LUs in a local SNA major node
- TRLEs in a TRL major node
- GROUP definitions, LINE definitions, and PORT connection network definitions in an Enterprise Extender XCA major node
- CDRSCs in a CDRSC major node
- APPL definitions in an APPL major node

Dynamic change of operands is also supported for PUs and LUs in switched major nodes.

There are several ways to perform dynamic reconfiguration. See [“Dynamic reconfiguration and dynamic change requirements” on page 175](#) to determine which method is valid for a given resource:

- Add, delete, or move resources in the VTAMLST definition file and issue a VARY ACT,UPDATE=ALL command. Changes made in this way are permanent changes, and do not require the major node to be deactivated.

See [“Using the VARY ACT,UPDATE technique” on page 176](#) for additional information.

- Code a DR file to add, delete, and move resources and activate it with the VARY DRDS command. Changes made in this way are temporary changes.

See [“Using the VARY DRDS technique” on page 178](#) for additional information.

- Use the MODIFY DR command to move or delete resources. Changes made in this way are temporary changes.

See [“Using the MODIFY DR technique” on page 180](#) for additional information.

Dynamic reconfiguration and dynamic change requirements

Table 10 on page 175 shows which techniques of dynamic reconfiguration and dynamic change of operands are valid for a particular operation on a resource in a particular major node.

<i>Table 10. Dynamic reconfiguration operations for valid major nodes</i>						
Operation/resource	Local SNA	Switched	TRL	CDRSC	EE	APPL
Add LU	2	1	-	-	-	-
Delete LU	3	1	-	-	-	-
Move LU	1	-	-	-	-	-
Change LU operands	1	1	-	-	-	-
Add PU	-	1	-	-	-	-
Delete PU	-	1	-	-	-	-
Move PU	-	-	-	-	-	-
Change PU operands	-	1	-	-	-	-
Add PATH	-	1	-	-	-	-
Delete PATH	-	1	-	-	-	-
Change PATH operands	-	1	-	-	-	-
Add TRLE	-	-	1	-	-	-
Delete TRLE	-	-	1	-	-	-
Change TRLE operands	-	-	1	-	-	-
Add CDRSC	-	-	-	1	-	-
Delete CDRSC	-	-	-	1	-	-
Change CDRSC operands	-	-	-	1	-	-
Add GROUP	-	-	-	-	1	-
Delete GROUP	-	-	-	-	1	-
Change GROUP operands	-	-	-	-	1	-
Add LINE	-	-	-	-	1	-
Delete LINE	-	-	-	-	1	-
Change LINE operands	-	-	-	-	1	-

Table 10. Dynamic reconfiguration operations for valid major nodes (continued)						
Operation/resource	Local SNA	Switched	TRL	CDRSC	EE	APPL
Change PORT connection network values	-	-	-	-	1	-
Add APPL	-	-	-	-	-	1
Delete APPL	-	-	-	-	-	1
Change APPL operands	-	-	-	-	-	1

Legend:

- 1** V ACT,UPDATE only
- 2** Either V DRDS or V ACT,UPDATE
- 3** Either MODIFY DR, V DRDS, or V ACT,UPDATE
- 4** Implicit dynamic reconfiguration

Dynamic reconfiguration is not supported on a line that has type 4 or type 5 physical units on it (with the exception of type 2.1 PUs).

Using the VARY ACT,UPDATE technique

You can use the VARY ACT,UPDATE technique to dynamically change certain operand values on definition statements without deactivating the major node. Only the resources that you are changing must be inactive.

To change the values, edit the existing definitions in the VTAMLST file and then issue a VARY ACT, UPDATE=ALL command.

Some VTAM-only operands can be changed, depending on the state of the resource. All other operands can be changed only by deleting the resource and adding the resource. For more information about which operands support dynamic change, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

To use the VARY ACT,UPDATE technique for dynamic reconfiguration, change the VTAMLST source file for a major node and issue a VARY ACT,UPDATE=IMPLICIT|ADD|ALL command. The command causes VTAM to update the actual configuration to match the new definition in the VTAMLST source file. The major node can be active for ADD or ALL. Only the individual resources being added, deleted, or moved must be inactive.

Restrictions for changing PU and LU resources: When using the VARY ACT,UPDATE=ALL command to dynamically change PU and LU resources, the following additional restrictions apply:

- The PUTYPE value cannot be changed. Also, you cannot add SDDL support to a PU.
- You can convert an APPN PU to a LEN PU, or a LEN PU to an APPN PU, by changing the value of the CONNTYPE operand. However, you cannot convert an APPN PU to a LEN PU if one of the following is true:
 - You coded the same NETID/CPNAME value on multiple PUs.
 - You coded an ADJCP definition statement for a PU to represent a partner CP and you coded the same NETID/CPNAME value on that PU.

Restrictions for changing Enterprise Extender resources: When using the VARY ACT,UPDATE=ALL command to dynamically change Enterprise Extender resources, the following additional restrictions apply:

- Changes to Enterprise Extender groups require the deactivation of the group and all subordinate resources within the group. This is done by using the VARY INACT,ID=*groupname* command.
- Changes to port connection network values require the deactivation of the group identified by the VNGROUP operand.

Restrictions for changing model CDRSCs: When using the VARY ACT,UPDATE=ALL command to dynamically change model CDRSCs in a CDRSC major node, the following additional restrictions apply:

- Existing clone CDRSCs are unaffected by changes to the model CDRSC.
- A model CDRSC cannot be deleted while it has clone CDRSCs.
- A conventional CDRSC cannot be added when there is a clone CDRSC of the same name. First, use the VARY INACT,DELETE=YES command to remove the clone CDRSC and then add the conventional CDRSC using the VARY ACT,UPDATE=ALL command.
- V ACT,UPDATE=ALL cannot be used to delete clone CDRSCs.

Restrictions for changing QDIO TRLE definitions: Stop all TCPIP devices that are using the TRLE and 2 minutes must pass before the QDIO TRLE can be updated.

Restrictions for changing applications: When using the VARY ACT,UPDATE=ALL command to dynamically change applications in an application major node, the following additional restrictions apply:

- An application cannot be added when there is a clone application of the same name. First, initiate close processing for the clone application. This causes it to be deleted. Then add the application using VARY ACT,UPDATE=ALL.
- An application cannot be added to the major node unless it can be added to the host name space as a resource or a shadow resource.
- The VARY ACT,UPDATE=ALL command cannot be used to delete clone applications.
- A model application cannot be deleted if any clone application created using the model exists.

VARY ACT,UPDATE=IMPLICIT (the default) has the same effect as implicit dynamic reconfiguration. For more information about the VARY ACT command, see [z/OS Communications Server: SNA Operation](#).

With VARY ACT,UPDATE=ADD|ALL, dynamic reconfiguration changes stay in effect regardless of VTAM or NCP recycles (halt and restart) and upon takeover of an NCP. You do not need to redrive dynamic reconfiguration changes with VARY DRDS or MODIFY DR commands. It is recommended that you do *not* make changes to VTAMLST files during a VTAM or NCP recycle. Instead, make the changes dynamically while VTAM and NCP are running. To prepare for takeover of an NCP, distribute the new VTAMLST file to the backup host and issue the same VARY ACT,UPDATE=ADD|ALL command that you issued at the primary host.

Sifting of operands does take place on subordinate resources that have been added or moved by the VARY ACT,UPDATE method of dynamic reconfiguration.

Network management programs that rely on VTAMLST definitions for the network configuration have current information available to them about the configuration.

Notes: If you use the VARY ACT, the UPDATE=ALL command might override the following commands:

- The VARY LOGON command to establish an automatic logon specification that is not coded on the LOGAPPL operand of a definition statement
- The VARY NOLOGON command to delete an automatic logon specification that is coded on the LOGAPPL operand of a definition statement
- The MODIFY DEFAULTS, DLOGMOD command to change the value of the DLOGMOD operand of a definition statement
- The MODIFY TABLE command to change VTAM table associations

For example, the LOGAPPL operand in the VTAMLST definition overrides the VARY LOGON and VARY NOLOGON commands. Even if no LOGAPPL is coded in the VTAMLST definition, a null value for LOGAPPL overrides the VARY LOGON value. To avoid having the VARY LOGON command overridden, code the

LOGAPPL definition statement with the required value. To avoid having the VARY NOLOGON command overridden, delete the LOGAPPL operand from the definition statement.

Coding dynamic reconfiguration changes

To add resources to a network, code the new definition statements in the required location in the VTAMLST file for the major node. To delete resources, delete the appropriate definition statements from the VTAMLST file. To move resources, move the appropriate definition statements from one location to another location in the VTAMLST file.

You can move a physical unit (and its associated logical units) from a line under one major node to a line under a different major node under the same VTAM. To do this, delete the resources from one major node and add them to the other major node. You delete the resources with any method of dynamic reconfiguration. Make the additions in the VTAMLST file.

Some move or add operations can depend upon a delete operation completing first. For example, you could exchange resources between two major nodes at the same time (that is, move resources from the first major node to the second major node and also move resources from the second major node to the first). Or, you might delete a PU from a line and add a PU to the line with the same ADDR value as the deleted PU.

You can resolve this problem in either of these ways:

- Make deletions with any method of dynamic reconfiguration, and then make the additions by modifying VTAMLST and issuing the VARY ACT,UPDATE=ADD command.
- Process the source definitions twice (issue the same VARY ACT command again). The deletions will take place on the first run of the command, and the additions that failed on the first run will take place on the second run.

Errors in the VTAMLST definition file

If you attempt dynamic reconfiguration and the VTAMLST file contains errors, you might receive unexpected results. Errors in BUILD or PCCU statements or errors in required operands can cause the VARY ACT command to fail, in which case no dynamic reconfiguration takes place. Errors in definition statements can cause unplanned dynamic reconfiguration changes or unplanned dynamic changes in operand values. Depending on the severity of the error in the definition file, VTAM may skip over resources that have incorrect definitions and process the rest of the definition file. VTAM deletes resources that are skipped unless they are active.

Note: You should run the NCP/EP definition facility (NDF) or issue the VARY ACT,SCOPE=SYNTAX command to check for definition errors *before* you attempt dynamic reconfiguration changes.

Using the VARY DRDS technique

The VARY DRDS command activates a dynamic reconfiguration (DR) file.

Procedure

Take the following steps to use the VARY DRDS command:

1. Define the appropriate dynamic reconfiguration files.
2. Issue the VARY DRDS command to implement the changes.

Results

You can use a dynamic reconfiguration file to:

- Move a physical unit from one nonswitched line to another within the same NCP
- Add physical units to nonswitched lines, or add logical units to physical units under nonswitched lines
- Change the link station address of physical units under nonswitched lines

- Delete physical units from nonswitched lines and delete logical units from physical units under nonswitched lines
- Delete NCP frame-relay PUs

The NCP to be reconfigured must be active, and the resource to be deleted or moved must be inactive. Also, the NCP cannot recognize, move, or add requests unless the physical unit was active at some time before deactivating it to issue a VARY DRDS command on it.

Dynamic reconfiguration with a DR file should be used only as a temporary method; a new NCP should be generated to reflect the changes as soon as time permits. The NCP should be reloaded to reflect the permanent changes.

If you are coding a dynamic reconfiguration file:

1. The first definition statement in the file must be VBUILD TYPE=DR. Include ADD, DELETE, or MOVE statements for the changes you want to make. For each statement, provide VTAM definitions for the resources you want to change.

Code the proper value (LOCAL) for the DRTYPE operand on the ADD and DELETE definition statements to indicate that the dynamic reconfiguration operation is for a connection through the local peripheral node. The default (DRTYPE=NCP) does not work in this case.

2. File the dynamic reconfiguration file in the VTAM definition library.
3. Make sure that you have deactivated all resources that must be inactive.
4. Issue the VARY DRDS command to invoke dynamic reconfiguration using the dynamic reconfiguration file.
5. Issue the VARY ACT command to activate newly added resources and resources you deactivated for the reconfiguration.

To automatically activate a resource added by dynamic reconfiguration, the status of the higher-level node should be ACTIVE, and ISTATUS=ACTIVE should be specified on the newly added resource.

To automatically activate a moved physical unit and its logical units that have ISTATUS=ACTIVE coded, specify ACTIVATE=YES on the PU definition statement following a MOVE statement.

6. Check the results of the dynamic reconfiguration by issuing a DISPLAY command for the affected resources.

If the NODELST start option is in effect when you perform a dynamic reconfiguration, VTAM records the name of each dynamic reconfiguration file in the NODELST file. When you restart VTAM after a halt or a failure, you can automatically reapply the changes in the dynamic reconfiguration files by specifying the name of the NODELST file on the CONFIG start option.

If a physical unit was dynamically moved before VTAM failed, it cannot be moved during a warm start. The activation of resources precedes the processing of dynamic reconfiguration files during warm start. Because dynamic reconfiguration depends upon resources being inactive, dynamic reconfiguration move operations are not performed. Deactivate the PU and issue the VARY DRDS command again.

Note: VARY ACT,UPDATE enables resources to be dynamically reconfigured and activated on warm start.

Adding resources

To add resources, code the PU and LU definition statements as you would when defining these resources. If you do not code the MAXDATA operand on the PU definition statement, a value of 261 bytes is used for a PU type 1 and a value of 265 bytes is used for a PU type 2. Code the operands on the PU and LU definition statements in the same format as they appear in the NCP major node. For example, ensure that PASSLIM and MAXOUT are coded correctly. See the *NCP, SSP, and EP Resource Definition Reference* for the coding format of NCP operands.

In a DR file, explicitly code all operands from the GROUP, LINE, or PU to which you are adding the PU or LU. VTAM does not process most NCP-only operands and might issue a warning message, but the GP3174 and RETRIES=(t,n) NCP operands can be used when adding PUs on a nonswitched SDLC link.

Sifting takes place within the hierarchy of minor nodes being added dynamically. Values for operands coded in the original hierarchy above the hierarchy being added with the DRDS deck do not sift down to the added resources. Therefore, if you want the values coded in the original hierarchy, code them explicitly in the DR file for the resources you are adding.

The most efficient way to add a physical unit and its associated logical units with a DR file is to specify them on the same ADD statement. It is not necessary to use separate ADD statements for the physical unit and each logical unit.

Deleting resources

To delete resources, code only the PU or LU names in the DR file. You do not have to code individual PU or LU operands. When you delete a PU, you also delete its associated LUs. Do not code the associated LU definition statements.

Moving resources

To move resources using a DR file, code only the names of the PUs you want to move. Do not code associated LUs. VTAM automatically includes the associated LUs in the move. For each PU, you can optionally specify a new SDLC link station address and whether you want the PU and its LUs automatically activated after the move.

To change the SDLC link station address for a physical unit on a certain line, use the MOVE statement and specify the new address for the ADDR operand. Keep the FROM and TO line the same.

The MOVE operation through a DR file is only for PUs (and their associated LUs). To move an individual LU to another PU, use DELETE and then a subsequent ADD. (For efficient storage management, code all DELETES before their related ADDs.)

Note: You can move a physical unit only under the same NCP major node. The physical unit cannot be moved from a line attached on one NCP to a line associated with another NCP.

Using the MODIFY DR technique

Using the MODIFY DR command, you can move a PU (and its LUs) from one line to another line, delete a PU (and its LUs) from a line, or delete an LU from a PU. You can change the SDLC address for a PU by requesting a move in which the new line and old line are the same, but the address is different.

The NCP to be reconfigured must be active, and the resource to be deleted or moved must be inactive.

Dynamic reconfiguration with the MODIFY DR command should be used only as a temporary method; a new NCP must be generated to reflect the changes as soon as possible. The NCP must be reloaded to reflect the permanent changes.

Dynamic reconfiguration changes made by MODIFY DR are not recorded in NODELST. Therefore, they are not automatically applied and can cause errors because they are not used during a warm start.

Dynamic reconfiguration of independent LUs

You can dynamically add or move independent LUs with the VARY DRDS command or the VARY ACT command with the UPDATE operand. You can also dynamically change operand values in an independent LU with the VARY ACT,UPDATE=ALL command. For a dynamic reconfiguration add or move, or for dynamically changing operand values, the processing takes place as shown in [Table 11 on page 180](#). For further information, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Table 11. Rules for multiple definition of resources			
Existing resource	Input resource	Major node after integration	Notes
Dynamic CDRSC	DR-defined independent LU	ISTPDILU	“1” on page 181 , “2” on page 181

Table 11. Rules for multiple definition of resources (continued)

Existing resource	Input resource	Major node after integration	Notes
Predefined independent LU	DR-defined independent LU	ISTPDILU	“1” on page 181 , “2” on page 181
Predefined CDRSC	DR-defined independent LU	ISTPDILU	“1” on page 181 , “3” on page 181

Notes:

1. If the existing resource has a known real NETID, it must match the NETID of the input resource. If the real NETID of the existing resource is not known, the NETID is inherited from the input resource.
2. All operands from the input resource apply to the existing resource.
3. Only LU operands from the input resource apply to the existing resource.

Chapter 9. Defining peripheral nodes

Peripheral nodes use local addresses for routing and require boundary function assistance from VTAM or the NCP. Logical units (LUs) are the ports through which users access the network. The type 1 and type 2 peripheral node architecture supports dependent LUs only. The type 2.1 peripheral node architecture supports independent and dependent LUs.

Defining type 2.1 peripheral nodes

For switched connections, PU definitions for type 2.1 peripheral nodes can be created dynamically on the VTAM host accepting the incoming call. For details, see [Chapter 8, “Defining resources dynamically,” on page 167](#).

If you code PU definition statements for type 2.1 peripheral nodes, consider the following.

- **Type 2.1 characteristics**

Type 2.1 peripheral nodes can communicate with each other, and with application programs in VTAM. Communication between type 2.1 peripheral nodes can occur through the main communication path among the subarea nodes in the network. Type 2.1 peripheral nodes support link-level role negotiation, eliminate hierarchical control (no ACTPU or ACTLU unless PU requests), require fewer flows to establish sessions, and use dynamically assigned session identifiers rather than preassigned addresses, all of which require fewer system definitions.

VTAM performs the directory services function for the subarea network and any APPN end nodes it serves. It locates the other session partner through a type 2.1 logical unit session request. VTAM is also involved in the session setup process when the type 2.1 session request traverses the subarea network, or when it is acting as a network node server for an APPN end node. VTAM or VTAM and NCP together appear and operate as a peer node to any type 2.1 peripheral node that attaches to the subarea network.

- **CPNAME operand**

For a type 2.1 peripheral node on a switched or LAN connection, VTAM requires the coding of either the CPNAME operand or both the IDBLK and IDNUM operands on the PU definition statement in a switched major node. You can code all three. During activation of the node, VTAM uses the NETID (if specified) and the control point name (CPNAME) to find the PU and LU definition statements associated with the node. If CPNAME is not in the exchange ID (XID) or if CPNAME is not coded in VTAM, VTAM uses IDBLK and IDNUM instead. (To have VTAM first use IDBLK and IDNUM, use the SWNORDER=STATNID start option or specify the SWNORDER=STATNID operand on the GROUP or LINE definition statement.) VTAM also uses the control point name to locate switched node definitions, even if the boundary function does not provide type 2.1 peripheral node support.

- **XID services**

For a type 2.1 peripheral node attached to the communication adapter using an SDLC link, the MODE operand on the LINE definition statement can be used to indicate whether VTAM performs the functions of the primary (PRI) or the secondary (SEC) link station. VTAM uses this PRI or SEC indication to determine its link-level role during the exchange identification (XID) processing for the connection. Because VTAM does not negotiate the primary or secondary role and uses the MODE specification for the XID process, you need to know the link-level role capability of the other type 2.1 connection. For example, if you specify VTAM as primary (MODE=PRI) and the other peripheral node is also primary nonnegotiable, the connection fails.

If the SDLC line used is a switched link, the XMITDLY operand of the LINE definition statement can be used to specify the amount of time that VTAM should delay before performing XID processing after answering the incoming call. This delay allows the calling peripheral node to initiate XID transmission first and, therefore, to avoid VTAM initiating this process before the type 2.1 peripheral node is ready.

The XID=YES operand coded on the PU definition statement implies that the contact procedure used is for a type 2.1 peripheral node. VTAM determines whether the peripheral node is a type 2.1 or a type 2 by checking the XID value received from the node either during activation of the node or at CONTACT time for switched nodes.

- XCA local area networks

Although VTAM can support type 2.1 peripheral nodes that are channel-attached through the external communication adapter (XCA), review the connectivity capabilities of each particular device to ensure it can be channel-attached as a type 2.1 peripheral node.

- Defining a type 2.1 channel

If you are connecting a channel-attached type 2.1 PU to your host, code a VBUILD definition statement with TYPE=LOCAL specified.

```
LSNA3    VBUILD TYPE=LOCAL
LSNA3PU  PU      CUADDR=051,
                MAXBFRU=15,
                PUTYPE=2,
                SSCPFM=USSSCS,
                VPACING=0,
                XID=YES
```

Nonnative network type 2.1 connections

You can connect to type 2.1 PUs that are not in your network but instead are in a nonnative network. You can use the XNETALS start option or code the XNETALS operand on the GROUP, LINE, or PU statement, to network qualify any LUs that you connect to in a nonnative network.

If XNETALS is resolved to YES at the PU level, VTAM uses the NETID of the PU (rather than the VTAM network identifier) as the network qualifier of the LU name. For example, in [Figure 51 on page 184](#), XNETALS is resolved to YES at the PU level and the LU in the attached type 2.1 peripheral node is recognized by VTAM as AS400.LU even though the network identifier of the subarea network is SANETX.




Figure 51. Nonnative network type 2.1 connection

Coding the NETID operand on the PU definition statement is optional (VTAM determines the network ID if not specified). However, if you specify XNETALS=YES, code the NETID operand if you want VTAM to perform dial-out operations to establish a session.

If XNETALS is resolved to NO at the PU level, and you do not code the NETID operand on the PU definition statement, VTAM assumes all connections are native unless the PU is a dependent LU server (DLUS) PU or a dependent LU requester (DLUR) PU. For information about DLUS and DLUR, see [“Dependent LU server” on page 405](#).

If you code the NETID operand on the PU definition statement and XNETALS=NO is specified for the PU, the PU NETID must match the VTAM NETID.

Attaching peripheral nodes to VTAM

Peripheral nodes can be attached to VTAM with the following channels:

- Local non-SNA
- Local SNA
- Loop adapter attached
- External communication adapter

Local non-SNA connection

Each local non-SNA major node defines a set of channel-attached (local) non-SNA terminals. Each minor node represents a non-SNA terminal (such as a 3277). You should deactivate a local non-SNA major node before you take its terminals offline.

Procedure

Take the following steps to define a local non-SNA peripheral node:

1. Code an LBUILD definition statement followed by a collection of VTAM definition statements. These definition statements should define each channel-attached non-SNA terminal as part of a logical set (group) of channel-attached non-SNA terminals.

Note:

- You do not need to code a definition statement for the non-SNA cluster controller (3272 or compatible device) to which the terminal is attached. Non-SNA terminals connected to a single cluster controller do not have to be defined in a single major node. Different terminals on the same controller can be defined to VTAM in different local non-SNA major nodes.
 - You cannot specify the channel device name when you are activating the logical unit for a channel-attached non-SNA device; the channel device name must be specified on the LOCAL definition statement for the device.
 - Code one LBUILD definition statement for each logical group (major node) of channel-attached non-SNA terminals.
 - Do not code a PU definition statement. The physical unit for a channel-attached non-SNA device is the VTAM host physical unit (ISTPUS), which is always active while VTAM is running. Therefore, only the local non-SNA major node and its logical units (representing the channel-attached devices) need to be defined and activated.
2. Code a LOCAL definition statement for each terminal (minor node) in the group. One or more LOCAL definition statements can be grouped with an LBUILD definition statement.

Results

Following is an example of definition statements for a local non-SNA major node:

LC3270NS LBUILD	LOCAL NON-SNA ATTACHMENT
LOC3277 LOCAL CUADDR=00A,	CHANNEL UNIT ADDRESS
TERM=3277,	TERMINAL TYPE
:	
DLOGMOD=S3270,	DEFAULT LOGON MODE TABLE ENTRY
LOGAPPL=A50ACCTS	AUTOMATIC LOGON APPLICATION

Note: A local non-SNA terminal should not be defined to and activated by VTAM if its channel unit address is defined as an MVS console and allocated to console services. Activating a local non-SNA terminal whose channel unit address is in use by console services can cause VTAM, console services, or both, to abend.

Local SNA connection

Each local SNA major node defines a set of channel-attached (local) SNA cluster controllers. Each minor node represents an individual physical or logical unit.

Channel-attached SNA devices can be either statically or dynamically defined. For an overview of defining local SNA devices dynamically, see [“Dynamic configuration of channel-attached devices” on page 170](#).

Defining a local SNA connection

You can define a local SNA device either statically or dynamically.

Procedure

Take the following steps to statically define a local SNA device:

1. Define a local SNA major node by coding a VBUILD definition statement for the major node and a separate PU or LU definition statement for each minor node.
2. Code a PU definition statement for each physical unit (such as a cluster controller) in the major node.

Notes:

- a. A physical unit and all its logical units must be defined within a single major node.
 - b. When activating a physical unit in a local SNA major node, supply the channel device name of the physical unit unless you have defined it in the PU definition statement. You can also specify the channel device name to override the value in the PU definition. To supply the channel device name, use the U operand of the VARY ACT command when activating the physical unit.
3. Code an LU definition statement for each logical unit placed after the associated PU definition statement. When you define a logical unit that logs on to VM, its name cannot be the same as any VM user ID in the system.

Sample local SNA connection

Figure 52 on page 186 is a sample configuration with local SNA devices.

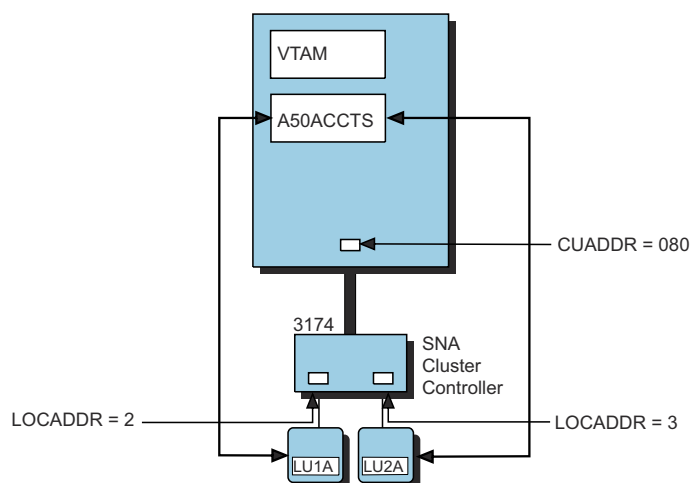


Figure 52. Local SNA devices

Following is an example of definition statements for [Figure 52 on page 186](#):

LOCALSNA	VBUILD	TYPE=LOCAL	LOCAL SNA ATTACHMENT
LPU1	PU	CUADDR=080,	CHANNEL UNIT ADDRESS
		...	
		MODETAB=LOG3174	LOGON MODE TABLE
*			
LU1A	LU	LOCADDR=2,	LOGICAL UNIT ADDRESS
		...	
		MDLTAB=MDLTAB1,	MODEL NAME TABLE
		MDLENT=ENTRY3,	MODEL NAME TABLE ENTRY
		LOGAPPL=A50ACCTS	AUTOMATIC LOGON APPLICATION
*			
LU2A	LU	LOCADDR=3,	
		...	
		MDLTAB=MDLTAB1,	MODEL NAME TABLE

MDLENT=ENTRY3,	MODEL NAME TABLE ENTRY
LOGAPPL=A50ACCTS	AUTOMATIC LOGON APPLICATION

If VTAM requires operator assistance to create a connection to a physical unit in a local SNA major node, VTAM displays a message indicating that a connection request was denied because the physical unit is offline. The connection request remains pending and the operator should either allow the connection request to complete by making the device available, or disallow the connection request by entering a VARY INOP command for the physical unit.

Enterprise Systems Connection (ESCON) channel attachment (ESA systems only)

The Enterprise Systems Connection (ESCON) channel is a high-bandwidth host attachment facility that can be used to connect SNA and non-SNA 3174 controllers, 3172 Nways interconnect controllers, 3746-900 controllers, and channel-attached hosts to VTAM running on MVS.

Loop-adapter-attached connection

Loop-adapter-attached devices are defined in a local SNA major node, and they communicate with VTAM application programs as if they were channel-attached SNA physical units and logical units.

Chapter 10. Defining LUs

Logical units (LUs) are the ports through which users access the network. Type 1 and type 2 peripheral node architecture support dependent LUs only. Type 2.1 peripheral node architecture supports independent and dependent LUs.

Independent LUs

Independent LUs are represented to VTAM as cross-domain resources (CDRSCs). Independent LUs can be coded with LOCADDR=0 in major nodes, coded as CDRSCs, or dynamically defined, but in all cases, VTAM represents them as CDRSCs. The independent LUs can access the network from multiple adjacent link stations, which can be in any domain. Even if you predefine independent LUs in device-type major nodes (NCP, local SNA, and switched major nodes), VTAM converts these definitions to CDRSCs when the major node is activated. Because independent LUs are represented as CDRSCs, when you use a DISPLAY command on an independent LU, the VTAM message responds with TYPE=CDRSC. However, message IST1131I shows a device type of INDEPENDENT LU / CDRSC.

Characteristics of independent LUs

Following are some characteristics of independent LUs.

Adjacent link stations

Adjacent link stations are type 2.1 physical units (or type 4 or type 5 PUs that appear as type 2.1 PUs) through which independent LUs can access the network. There is no hierarchical relationship between an adjacent link station (PU) and an independent LU. For more information about adjacent link stations and how independent LUs can use them to access the network, see [“Using CDRSC definition statements for independent LUs” on page 193](#), and [“Multiple connections between a type 2.1 node and a subarea node” on page 194](#).

Session capabilities

The session capabilities of an independent LU are significantly greater than those of dependent LUs. Independent LUs do not require SSCP-PU or SSCP-LU sessions. An independent LU can act as a primary LU or a secondary LU to initiate sessions. However, an independent LU must be the secondary logical unit if it is to be the destination LU (DLU) of a session.

Extended BIND

Independent LUs support the extended BIND, which contains information necessary to activate a session, to make use of features such as adaptive session pacing, and to identify a particular session. Extended BINDs always contain a fully qualified procedure-correlation identifier (a PCID qualified by the name and network identifier of the PCID generator). A BIND without this information is a nonextended BIND. Extended BINDs can contain network-qualified names if the subarea components in the session establishment path support them. An extended BIND is sent only if the boundary function (if applicable) and the destination LU support extended BIND.

Multiple sessions

An independent LU can also have multiple sessions with a single LU (parallel sessions) or with several different LUs. The independent LU can assume the role of primary or secondary for any one or all of the sessions.

If one or more sessions are required between LUs for application transaction program communication, the independent logical unit establishes the necessary session connections on behalf of those application transaction programs.

The BSBUF buffer pool is used to provide common storage for boundary type 2.1, type 2, and type 1 peripheral-node-session control blocks. Depending on how you code the buffer pool start option, the BSBUF buffer pool can expand to support more sessions and contract as less session control blocks are needed. Because independent LUs can have multiple sessions (including parallel sessions), you should determine the average and maximum number of sessions associated with independent LUs that are directly attached to that VTAM subarea node. You can then specify a value that prevents excessive buffer expansion and contraction.

Addressing

An independent LU in a type 2.1 peripheral node does not have a predefined address. Although VTAM definition statements associates an independent LU with a network accessible unit, the address of that unit does not have to be coordinated with VTAM coding.

Alias name translation and independent LUs

The following restrictions apply to alias name translation and the independent LUs:

- If the PLU name received in the BIND from the LU is network-qualified, the SLU name must also be network-qualified.
- If the names are network qualified, they must be the real names (unless they are USERVAR names), or the session setup fails.
- Concurrent sessions to destination resources with the same name in different networks cannot be established when using alias-name translation. Concurrent sessions to destination resources with the same name in different networks can be established if the independent LU host is using NQNM= NQNAME to define the destination resources. For more information about NQNM= NQNAME, see the [z/OS Communications Server: SNA Resource Definition Reference](#).
- If the SLU name received in the BIND is network qualified, and the SLU is from a different network than the PLU, the SLU alias name in the PLU network will be the same as the SLU real name.

Dynamic reconfiguration and independent LUs

You can create and alter CDRSC major nodes as needed to effect the same level of support as dynamic reconfiguration provides. Because dynamic reconfiguration has meaning only if the independent LUs are predefined using LU definitions, it is no longer needed when you convert your independent LU definitions to CDRSCs.

Note: The ISTDILU major node contains CDRSC definitions for independent LUs that are converted from LU definition statements by VTAM. The ISTDILU major node is automatically activated by VTAM on initialization and cannot be activated or deactivated by operator commands.

When processing a dynamic reconfiguration definition that contains an independent LU definition to be added, VTAM converts that independent LU to a CDRSC and places it in the ISTDILU major node. The PU named on the TO operand becomes an entry in the CDRSC default adjacent link station list.

If you specify a PU on the FROM operand of the dynamic reconfiguration DELETE command on an independent LU, VTAM deletes the network address associated with that independent LU that was predefined for that independent LU. Additionally, the PU is removed from the default adjacent link-station list if present.

To use the dynamic reconfiguration DELETE command on an independent LU, the independent LU must not have any sessions over the PU named on the FROM operand.

You cannot use the MODIFY DR command for independent LUs.

Defining independent LUs

You can define independent LUs in the following ways:

- You can have your independent LUs defined dynamically by VTAM during session establishment. See [“Dynamic definition of independent LUs” on page 191](#).

- You can define CDRSC definition statements (or model CDRSC definition statements) to represent your independent LUs. See [“Using CDRSC definition statements for independent LUs”](#) on page 193.
- You can use standard LU definition statements with LOCADDR=0 coded. These LU definitions are converted into CDRSC definitions by VTAM when the major node is activated. This method is not recommended other than as a migration tool. Using this method, you cannot warm-start VTAM and have your independent LUs return to a before-failure condition. Independent LUs that are defined this way are converted to CDRSCs by VTAM and placed in the ISTPDILU major node. If a new resource is detected that matches an existing CDRSC, the new resource definition is integrated with the existing definition, and the resulting resource resides in the ISTPDILU major node. See [“Automatic conversion of independent LU definition statements”](#) on page 193.

For information about how VTAM handles multiple definitions of one independent LU resource, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Migration: It is recommended that you remove predefined independent LU definitions from your device-type major nodes and either let VTAM dynamically define CDRSC definitions for the independent LUs on a session-request basis or create CDRSC major nodes. If you decide to code CDRSC major nodes for your independent LUs and you previously had them defined to the NCP, you should regenerate your NCP to avoid any address mismatch problems.

One of the advantages of removing existing LU definition statements for independent LUs (and using either dynamic definition of independent LUs or coding CDRSC definition statements) is that predefined, NCP generated addresses are not used by the independent LU. Addresses are requested as needed and are returned when no longer needed.

You can use multiple connections and the dynamic selection of session connections without dynamically defining your independent LUs as CDRSCs.

Note: If you have session takeover operations that rely on VTAM creating substitute LUs, these takeover operations can fail.

Dynamic definition of independent LUs

Independent LUs attaching to VTAM through a type 2.1 PU can be dynamically defined as cross-domain resources. No system definition of the independent LU is required.

If the independent LU initiates the session, the adjacent link station (ALS), through which the independent LU contacts the network, does not have to be predefined to VTAM. For example, if the independent LU is attached through Enterprise Extender, the adjacent link station (switched PU) would be defined with a PATH definition which specifies either the IPADDR or HOSTNAME of the partner host.

If the independent LU is the destination LU unit rather than the originating LU, and the destination LU is located on a LEN node, then you can dynamically define the independent LU using the adjacent link station (ALS) selection function of the session management exit routine. (An independent logical unit that is the destination LU on an APPN node has no such restriction.) Using this exit routine, you can provide VTAM with a method of contacting the independent LU. For information about the ALS selection function, see [“Dynamic selection of session connections”](#) on page 195. For information about writing a session management exit routine, see [z/OS Communications Server: SNA Customization](#).

Note: The ALS selection function of the session management exit routine is not required if the destination LU is a LEN CP independent LU. Use the CPCDRSC=YES start option for this situation. If a session request is made to a LEN CP independent LU and CPCDRSC=YES is coded, VTAM builds a dynamic CDRSC to represent the LEN CP independent LU. As a result, other network resources will be able to establish sessions with the LEN CP independent LU before it has attempted to establish any sessions.

To have your independent LUs dynamically defined, perform the following steps:

1. Code DYNLU=YES in the start list, or use it on the START command when starting VTAM; code DYNLU=YES on the adjacent CP (ADJCP) definition statement for the adjacent APPN node; or code DYNLU=YES on all PU definition statements that represent links to the adjacent node. Coding DYNLU=YES enables VTAM to learn about the independent LU during session activation and dynamically define representations of the LU. If you code DYNLU=YES when starting VTAM (either in

the start list or on the START command), you can restrict dynamically defined LUs from being used to connect to an adjacent node by coding DYNLU=NO on the ADJCP statement for the adjacent node or by coding DYNLU=NO on all PU definition statements that represent links to the adjacent node.

Results: The DYNLU value is associated with an adjacent node when the first link to that adjacent node is activated. When the DYNLU value is associated with the adjacent node, that value will be propagated to all other links when they are activated. See [z/OS Communications Server: SNA Resource Definition Reference](#) for further information about determining the source of the DYNLU value assigned to an adjacent CP and attached resources.

2. For virtual route transmission groups (VRTGs), the default for the CDRDYN start option is YES, enabling dynamic CDRSC definition for the host node. The CDRDYN operand can also be specified on the host CDRM definition statement, but it will be overridden by the CDRDYN start option. You must also code CDRSC=OPT on the CDRM definition statement for an adjacent node to allow dynamically defined CDRSCs for sessions with that node. If you code CDRDYN=YES on the host CDRM definition statement, or allow the CDRDYN start option to default to YES, you can restrict dynamically defined CDRSCs from being used for sessions with an adjacent node by coding CDRSC=REQ (or allow it to default) on the CDRM definition statement for the adjacent node.

Results: The DYNLU value is set for an adjacent APPN node from the CDRSC keyword associated with the CDRM for that node when the first link that is activated to the adjacent node represents a virtual route transmission group (VRTG).

3. Set the CDRSCTI start option if you want to change the amount of time that the dynamic definitions of independent LUs are retained after session termination (default is 8 minutes).
4. Code the CPCDRSC=YES start option in a host that is connected to a LEN CP to allow dynamic definition of the LEN CP independent LU when it is the destination LU. DYNLU still controls dynamic definition of all independent LUs (including LEN CPs) when they are the originating LU.

When an independent LU is dynamically created by VTAM, the following defaults are used:

- ENCR=NONE
- RESSCB=0
- EAS=256
- MODETAB=value specified on the DYNMODTB start option, if coded (if the DYNMODTB start option is not coded, ISTINCLM is used)
- DLOGMOD=value specified on the DYNDLGMD start option, if coded

The PACING and VPACING operands specified on the PU definition statement associated with the session are sifted down to the CDRSC. The default logon mode table for dynamic CDRSCs is specified using the DYNMODTB start option or the DYNMODTB operand on the MODIFY VTAMOPTS command. You can also replace the existing table that is being used by dynamic CDRSCs by using the MODIFY TABLE,OPTION=LOAD command. If a default logon mode table for dynamic CDRSCs is not specified, ISTINCLM is used. The logon mode for each session is extracted from the BIND request that is passed to VTAM.

If the adjacent link station that the independent LU is using is being dynamically defined using dynamic switched definitions, you need to ensure that you have an entry for the mode you want to use in the logon mode table for the resource. If the mode entry is not found in the logon mode table for the resource or in the default logon mode table, the session fails. You can specify a default logon mode entry for dynamic CDRSCs using the DYNDLGMD start option or the DYNDLGMD operand on the MODIFY VTAMOPTS command. You can also override this value for a specific dynamic CDRSC using the MODIFY DEFAULTS or MODIFY RESOURCE commands.

Note: Using the MODIFY VTAMOPTS command to change the value of the DYNMODTB or DYNDLGMD start options does not affect dynamic CDRSCs that have already been created.

Dynamic CDRSCs are created and maintained in the ISTCDRDY major node. ISTCDRDY is activated automatically during VTAM initialization, and deactivated automatically during VTAM termination. The operator can deactivate ISTCDRDY with the VARY INACT command, in which case all dynamically defined CDRSCs are also deactivated and the dynamic CDRSC function is disabled. In addition, all sessions

involving dynamically created CDRSCs (including CP-CP sessions with this host, if the partner CP was dynamically defined) are terminated.

For information about how VTAM handles multiple definitions of one independent LU resource, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

You can use dynamic definition of independent LU with other functions, such as multiple connections to a subarea, casual connection, or APPN.

Using CDRSC definition statements for independent LUs

You can code CDRSC definition statements (or model CDRSC definition statements) for your independent LUs and specify the adjacent link stations (PUs) that VTAM uses to contact the independent LU. You can specify the adjacent link stations either by using the ALSLIST operand on the CDRSC definition statement, or by using the adjacent link station selection function of the session management exit routine. For more information about using the exit routine, see [“Dynamic selection of session connections” on page 195](#). For more information about defining model CDRSCs, see [“Model definition of cross-domain resources” on page 428](#).

The following sample shows an independent LU defined using a CDRSC definition statement:

```
CDRSC      VBUILD TYPE=CDRSC
VRCDRSC    CDRSC  ALSLIST=(SWPUAI06,SWPUA107)  ADJACENT LINK STATION NAMES
                                           USE ALS LIST ONLY
                                           ALSREQ=YES,
                                           LOGAPPL=APPL1,
                                           DLOGMOD=DSILGMOD,
                                           RESSCB=3
                                           AUTOMATIC LOGON
                                           LOGON MODE TABLE ENTRY
                                           SESSION CONTROL BLOCKS
```

ALSLIST=*puname* indicates to VTAM that this resource is an independent LU defined as a CDRSC. The ALSLIST operand can list one or more adjacent link stations that VTAM uses. The ALSLIST operand can also include an ISTAPNPU entry, a reserved keyword that represents any APPN capable PU. This entry can be used instead of listing each APPN adjacent link station. For information about controlling connection to independent LUs that are accessible over both LEN and APPN connections, see [“Establishing and controlling sessions” on page 412](#).

You can add and delete adjacent link stations in the list of adjacent link stations using the MODIFY ALSLIST command. For more information about this command, see [z/OS Communications Server: SNA Operation](#).

You can use the adjacent link station selection function of the session management exit routine to dynamically add other adjacent link stations to the list and to perform other adjacent link station selection functions. For details, see [z/OS Communications Server: SNA Customization](#).

You can use the ALSREQ operand or start option to control which adjacent link stations can be used for session traffic to independent LUs. The ALSREQ operand indicates whether this independent logical unit can establish sessions with an adjacent link station that is not defined by the ALSLIST operand. ALSREQ=YES means that a predefined adjacent link station is required for session establishment. If you code ALSREQ=NO, any adjacent link station can receive a session request for this LU.

Notes:

1. Even when ALSREQ=YES is specified, if any APPN adjacent link station name (or ISTAPNPU) exists in the ALS list for the CDRSC, any APPN adjacent link station can receive a session request for this LU.
2. Any operand normally coded on an LU definition statement can be coded on a CDRSC definition statement when defining an independent LU. These operands only have an effect when the CDRSC is an independent LU.

Automatic conversion of independent LU definition statements

The ISTDILU major node contains CDRSC definitions for independent LUs that are converted by VTAM from LU definition statements. In converting the independent LU to a CDRSC, VTAM removes the implied hierarchy of the independent LU to a type 2.1 physical unit.

The converted CDRSCs can be used with multiple connections between a type 2.1 node and a subarea node. For this reason, the individual CDRSC minor nodes that are converted from LU definitions are not freed when the NCP (or other major node) in which they were originally defined is deactivated.

However, you can delete individual CDRSC minor nodes without deleting the entire major node. You can use the DELETE operand on the VARY INACT command on predefined CDRSC minor nodes.

When the LU definition statement is converted, the resource is represented as a CDRSC. This means that if the OWNER operand is specified on the LINE definition statement where the independent LU is predefined, it has no effect on the independent LU being converted to a CDRSC. The independent LU is converted to a CDRSC, but the PU under which it is predefined cannot be used as a session path until it is acquired by that host.

Multiple connections between a type 2.1 node and a subarea node

An independent LU can connect to the subarea network through only a type 2.1 PU. A type 2.1 PU can be connected to the subarea network through a type 4 or type 5 PU. The type 2.1 PU connections are called adjacent link stations. An independent LU can connect to multiple adjacent link stations. Thus, the independent LU can have multiple, simultaneous connections to the VTAM host.

If one connection fails, only the sessions using that connection must be reestablished; the other sessions are not affected. The independent LU can then use an alternative adjacent link station for any subsequent sessions. You can restart any sessions terminated by the first adjacent link-station failure and connect the independent LU through another adjacent link station.

Some of the independent LU session paths might involve type 2.1 nodes connected to a VTAM other than the VTAM of the required session partner. For these sessions, the intermediate VTAM views the independent LU as a CDRSC. The CDRSC that represents the independent LU can also have other concurrent sessions that use PU type 2.1 nodes connected to the VTAM of the required session partner.

If an independent LU is connected to VTAM by only LEN links and two application programs try to contact it, VTAM provides the first available adjacent link station specified on the ALSLIST operand for both application programs. Both application programs would then have parallel sessions through the same adjacent link station. If you want the application programs to use different adjacent link stations, use the adjacent link station selection function of the session management exit routine. For information about coding this exit routine, see [z/OS Communications Server: SNA Customization](#).

Figure 53 on page 194 shows how an independent LU can connect to VTAM with multiple connections.




Figure 53. Independent LU with multiple connections to VTAM

Following are sample definition statements to define the configuration in Figure 53 on page 194. Note that you can support independent LU as either destination logical units, originating LUs, or both:

- For independent LUs as originating LUs, code DYNLU=YES in your start list. VTAM determines the adjacent link stations dynamically.
- For independent LUs as destination LUs, code CDRSC minor nodes in a CDRSC major node. If the independent LUs are connected to VTAM by only LEN connections, include the ALSLIST operand.

Start list (ATCSTRxx)

If the independent LUs are originating LUs, code the start option for dynamic definition of independent LUs (DYNLU=YES).

NCP major node

Code the following (do not code LU definition statements):

...	LUDRPOOL	NUMILU=2	NUMBER OF INDEPENDENT LUS
LGROUP	GROUP	DIAL=NO	LEASED LINE
LLINE	LINE		
...			
PU1	PU		
...			
SGROUP	GROUP	DIAL=YES	SWITCHED LINE
SLINE	LINE		
...			
PU2	PU		
...			

Switched major node

Code a switched major node:

VBUILD	TYPE=SWNET	
PUA	PU	ADDR=03,
		CPNAME=A500N,
		DISCNT=YES,
		PUTYPE=2
		STATION ADDRESS
		CONTROL POINT NAME (TYPE 2.1)
		DISCONNECT FACILITY
		PU TYPE 2.1

Note: You can code the DYNLU operand on the PU statement rather than in the start list.

CDRSC major node

If the independent LUs are destination LUs, define your independent LUs as CDRSCs:

VBUILD	TYPE=CDRSC	
ILU1	CDRSC	ALSLIST=(PU1,PU2),
		ALSREQ=YES
ILU2	CDRSC	ALSLIST=(PU1,PU2),
		ALSREQ=YES
		ADJACENT LINK STATION NAMES
		USE ALS LIST ONLY
		ADJACENT LINK STATION NAMES
		USE ALS LIST ONLY

Dynamic selection of session connections

If an independent LU has an existing leased or switched LEN connection (and no APPN connection), VTAM can dynamically select session paths for outbound sessions from the boundary function to the independent LU. VTAM attempts to use leased connections first, and then tries active switched connections. If the first adjacent link station in the adjacent link station list is not available, VTAM tries to use an alternate connection from a default adjacent link station list that VTAM creates. When the first adjacent link station (the primary connection) does become available, VTAM uses that adjacent link station for any new sessions. However, sessions already allocated to the alternate adjacent link station do not change.

To control selection of an alternate session path when an independent LU is connected to VTAM by only LEN links, use the ALS selection function of the session management exit routine. VTAM then establishes the path between the independent LU and the boundary function; this continues to be the path for the remainder of that session. You can select a different adjacent link station for each session.

Using the ALSREQ start option, you can limit the adjacent link stations that an independent LU can use.

You can code the session management exit routine to choose a PU from the default list or another PU by overriding the default PU list. For information about writing a session management exit routine, see [z/OS Communications Server: SNA Customization](#).

The default adjacent link station chosen must:

- Act as a type 2.1 PU
- Reside in a major node that supports independent LUs
- Be active, if it is nonswitched

- Be active, connectable, or pending dial if it is switched

Authorized transmission priority for LEN connections

Normally it is desirable for a session between two independent LUs through a subarea network to use the same transmission priority for both type 2.1 LEN connections (entry and exit). By default, when VTAM receives a BIND over a type 2.1 LEN connection, the transmission priority field received on the BFINIT is passed back on the BFCINIT and through the subarea network on the BIND. This allows this same transmission priority to be used for the session by the type 2.1 LEN connection on exit from the subarea network.

To have VTAM ignore the transmission priority specified on the BFINIT, use the AUTHLEN=NO start option or the AUTHLEN=NO operand on a particular PU definition statement (or on a GROUP or LINE statement to sift down to the PU). VTAM then sets the transmission priority field in the BFCINIT to B'01' (medium priority), regardless of the transmission priority specified on the initial BFINIT.

Restrictions on using independent LUs

Following are some restrictions that apply to independent LUs:

- An independent logical unit cannot be the destination LU (DLU) of a session, unless it is the secondary LU (SLU). An independent LU cannot function as the primary logical unit for an autologon (logappl) session. However, an independent LU can be the SLU of an autologon (logappl) session, provided that the independent LU resides on one of the following nodes:
 - A peripheral (type 2.1) node that is APPN-attached to the same APPN network as the VTAM from which the independent LU autologon session will be initiated
 - A peripheral (type 2.1) node that is LEN-attached directly to the VTAM (or an owned NCP) from which the independent LU autologon session will be initiated. Autologon is not supported for independent LUs residing on peripheral (type 2.1) nodes that are LEN-attached to non-VTAM/NCP nodes (for example, when independent LUs are migrated from being NCP-attached to being DLUR-attached). This is because the independent LU autologon session is initiated by VTAM when the adjacent link station to the LEN-attached node is activated; if a LEN-attached independent LU is not directly attached to VTAM (or an owned NCP), then VTAM has no knowledge of the activation of that link station, and therefore cannot reliably initiate the autologon session.
- An independent LU can initiate sessions as either a PLU or SLU.
- A dependent LU can have only one active LU-LU session at a time. Thus, there is no session with the controlling application program when the LU is in session with another application program. For independent LUs, this is not the case. VTAM initiates a session between an LU and its controlling application program, even though there are already active LU-LU sessions. Each time that a session with a PLU that is not the controlling application program ends, and the independent LU does not already have a session with the controlling application program, a session is established with the controlling application program. For more information about automatic logon with independent LUs, see [“Automatic logons” on page 209](#).
- The name interpretation table function is not supported for independent LU sessions.
- The SSCP-PU session is optional and can be requested by the type 2.1 node during the XID exchange. However, the SSCP-PU session is one of the mechanisms that the type 2.1 node uses to transport network management device information (alerts) to the SSCP for delivery to the NetView program. An alternative mechanism used to report alerts is through a network management focal point function.

Dependent logical units

A dependent LU must be activated, controlled, and owned by a VTAM SSCP. Ownership of dependent LUs is determined when the SSCP establishes an SSCP-LU session. Dependent LUs have the following characteristics:

- They generally have unique, fixed, preassigned local addresses within the PU that are specified on the NCP or VTAM definition statements using the LOCADDR operand of the LU definition statement.

With dynamic definition and with switched connections, the address is defined during the PU and LU activation process. For more information, see [“Defining dependent LUs dynamically” on page 197](#).

- They are subordinate to a PU. If defined manually, they must be defined under a PU.
- They require an SSCP-LU session and an SSCP-PU session to be active to VTAM; these sessions must be established before any LU-LU session can be activated.
- They normally act as SLU only.
- They can have only a single session. Even an LU that supports LU 6.2 protocols can have only one session if it is defined as a dependent LU.

Defining dependent LUs dynamically

You can enable VTAM to define dependent LUs dynamically. One way to do this is to use the IBM-supplied selection of definitions for dependent LUs (SDDL) exit routine. With this exit, VTAM can define LUs dynamically when it receives information from the PU that the device is powered on, rather than when the major node is activated. The dependent LUs must be attached to the host through PUs that support Reply/PSID NMVTs, request units that are used for dynamic definition of dependent logical units. For example, you can use a 3174 cluster controller because it supports these request units.

VTAM uses model LU definition statements to build LU definitions. Define model LU definition statements based on information the PU sends to VTAM about the logical unit device. The default information that the PU sends is the seven-character machine type and model number. When the PU is activated or when a device powers on, the PU sends a Reply/PSID NMVT request unit to VTAM, which can trigger VTAM to build or change LU definitions. This NMVT contains the local address of each LU, a power indicator (on or off), the machine type and model number of the device, and optionally other device-dependent information needed to define the LUs. VTAM uses this information to choose an appropriate model LU definition statement to build an LU definition.

LUGROUP operand

To enable this facility, code the LUGROUP operand on the PU definition statement for the PU, and code an LU group major node. The LUGROUP operand on the PU definition statement specifies the name of the model LU definition group that VTAM uses when dynamically defining LUs. The LU group major node contains the model definition statements. Dynamic definitions for LUs are built using the model LU definitions contained in this major node.

Note: Any operands that you code on the PU definition statement will not sift down to dynamically created LU definition statements.

After you enable this facility, VTAM uses the IBM-supplied selection of definitions for dependent LUs (SDDL) exit routine to generate LU names for the dynamically built LU definition statements. The SDDL exit routine controls the selection of the model LU and the LU name. If you prefer to use your own naming convention or method of selecting an LU model, you can write your own exit routine. For information about coding your own exit routine, see [z/OS Communications Server: SNA Customization](#).

If you use the IBM-supplied SDDL exit routine, code the LUSEED and LUGROUP operands on the PU definition statement. The LUSEED operand provides a pattern name that is used with the SDDL exit routine to create a name for the dynamically created LU definition statements.

[Figure 54 on page 198](#) illustrates VTAM dynamically defining dependent LUs attached through a 3174 cluster controller, and the steps listed after the table explain the process.




Figure 54. Definition building for dynamically defined dependent LUs

1. The LU at local address 003 powers on. The 3174 sends VTAM a Reply/PSID NMVT indicating a 3278 model 005 device has powered on at address 003.
2. VTAM uses the Reply/PSID NMVT to do the following actions:
 - a. Check the 3174 PU definition statement for a definition of an LU at address 003
 - b. Determine if the 3174 supports dynamic definition of dependent LUs (LUGROUP operand is coded), if the 3174 PU definition statement has no definition of an LU at address 003.
3. VTAM matches the LU group name specified in the 3174 PU definition statement with an LU group name in the LUGROUP definition statement.
4. The default SDDL exit routine uses the seven-character machine type and model number as the model LU name. VTAM serially checks the LU group for a model LU statement that matches the model name supplied by the SDDL exit routine. In this example, VTAM finds the second definition statement (3278@). The characteristics specified in that statement are used.
5. To create the LU name, the default SDDL exit routine uses the value on the LUSEED operand. In this example, the three # characters specified on the LUSEED operand are replaced with the numeric local address. Therefore, the name of the LU is LU003. For more options, see the [z/OS Communications Server: SNA Resource Definition Reference](#).
6. VTAM builds a dynamic definition for the LU. The dynamic definition is rebuilt if that device powers off and a new device powers on at the same local address. Also, the dynamic definition is rebuilt if the device powers off and powers back on after its LU group major node has been deactivated and then reactivated.

Chapter 11. Establishing and controlling SNA sessions

Logical units (LUs) are defined to VTAM to provide access services to network resources. To tailor sessions, VTAM uses the unformatted system service (USS), interpret, logon mode, and Class of Service tables. VTAM also uses the logon-interpret routines and the virtual route selection and session management exit routines to authorize, establish, or terminate sessions. You can replace or modify these tables and routines. This, along with operands that can be specified in VTAM definitions, allows VTAM to adapt to your needs. The [z/OS Communications Server: SNA Resource Definition Reference](#) describes how to replace or modify the tables used in session establishment and termination and how to specify the network definition operands that relate to sessions. [z/OS Communications Server: SNA Programming](#) describes how VTAM initiates, controls, and terminates sessions. [z/OS Communications Server: SNA Customization](#) describes the exit routines.

Sessions between LUs can be initiated by the following actions:

- Session-initiation requests from a logical unit
- Automatic logon
- The application program using API requests (for example, OPNDST, SIMLOGON, or REQSESS)
- The VTAM operator using, for example, VARY LOGON commands

Normally, sessions are terminated by one of the following actions:

- Session-termination requests from an LU
- The application program using API requests (for example, CLSDST or TERMSESS)
- The VTAM operator, using VARY resource deactivation commands

[z/OS Communications Server: SNA Programming](#) describes the use of VTAM application programming interface (API) macros to establish and terminate sessions.

The following list shows where to find more information about establishing and controlling SNA sessions.

- [“Multicultural support” on page 199](#)
- [“Model name table” on page 201](#)
- [“Associated LU table” on page 202](#)
- [“Selecting session parameters for the logon mode table” on page 203](#)
- [“Automatic logons” on page 209](#)
- [“Session management exits” on page 215](#)
- [“Session-level pacing” on page 217](#)
- [“Logon and logoff requests from dependent logical units” on page 232](#)
- [“Unformatted logon requests using mixed-case passwords” on page 233](#)

Multicultural support

With multicultural support you can tailor messages and commands for users. Using the MVS Message Service (MMS) or the LANGTAB USS tables, you can define USS messages in several different languages.

Also, by default, the language passed to a VTAM application program comes from the logon mode entry specified by the USS LOGMODE parameter. With multicultural support, the language passed to the application program can be specified by the user, independent of the LOGMODE parameter.

For further information about multicultural support, see [“Multicultural support for user USS messages and commands” on page 200](#) and [“Multicultural support for the language passed to application programs” on page 201](#).

Multicultural support for user USS messages and commands

Without multicultural support, USS messages (USSMSG00–USSMSG14) and commands (LOGON, LOGOFF, and IBMTEST) came from either a USS table that was statically defined to a logical unit, or the IBM-supplied default table, ISTINCDT. With multicultural support, users can choose a USS table by specifying the LANGTAB operand on any valid USS command. This enables users to enter USS commands and receive USS messages in the language of their choice. The USS table associated with the logical unit using the LANGTAB operand remains in effect until another valid USS command with the LANGTAB operand is processed.

You can also define USS messages using the MVS message service (MMS), which enables you to define each USS message in several different languages. Users can specify which language to use to retrieve USS messages from the MMS by using the LANG parameter on any valid USS command. The language for USS messages associated to the logical unit using the LANG parameter remains in effect until another valid USS command with the LANG parameter is processed.

For example, USSMSG10 can be defined in multiple languages, using the LANGTAB USS tables or MMS.

Remember that the language specified for the user USS commands (from the LANGTAB operand) and messages (from the LANG or LANGTAB operands) remain in effect for a logical unit (terminal) after the user logs off an application program. One logical unit (terminal) might be used by several users who speak different languages. Because LANG and LANGTAB can be entered with null values to cancel their effects, it might be wise to define LOGON commands with default null values for LANG and LANGTAB. For example, you can code the following as a default for the LOGON command:

```
LOGON    USSCMD    CMD=LOGON,FORMAT=PL1
         USSPARM    PARM=APPLID
         USSPARM    PARM=LANG,DEFAULT='('
         USSPARM    PARM=LOGMODE
         USSPARM    PARM=DATA
         EJECT
```

Note: The coding for LANGTAB is the same as shown in the example for LANG.

Defining USS tables for use with the LANGTAB operand

A character-coded command that follows the USS command syntax rules (see the [z/OS Communications Server: SNA Resource Definition Reference](#)) is referred to as a USS command. A session-level USS table can be used to convert the USS command into a field-formatted SNA request. If a character-coded command violates these syntax rules, an interpret table must be used for the conversion.

The following restrictions apply when defining a USS table used for LANGTAB:

- The name of the USS tables you code must match the name the user specifies on the LANGTAB operand.
- The name of the USS table cannot be ISTINCNO or ISTCFCMM.
- FORMAT=DYNAMIC must be coded on the USSTAB macro to use the table for LANGTAB.
- If TABLE is coded on the USSTAB macro, it is ignored because the input translation table cannot be specified by the user.
- USS message text must comply with the rules detailed in the [z/OS Communications Server: SNA Resource Definition Reference](#).

Defining USS messages to the MVS message service

To define USS messages to MMS:

1. Create message skeleton files for MMS. These files define the user messages (USSMSG00 – USSMSG14) in the languages required for your installation.
2. Define data sets to MVS.
3. Run the message compiler.

4. Update the parmlib.
5. Activate MMS.

For more information about message translation using MMS, see [Appendix G, “Message translation using the MVS Message Service,”](#) on page 609. For information about performing these functions, see [z/OS MVS Planning: Operations](#). For information about the format of message skeleton files used for language translation, see the [z/OS MVS Programming: Assembler Services Guide](#).

When defining a USS message in a message skeleton file, use the substitute tokens &1., &2., and &3. for the first, second, and third variable data fields for the message. These tokens correspond to the %(1), %(2), and %(3) syntax of specifying variable data that is used for the TEXT operand of the USSMSG macro instruction.

Multicultural support for the language passed to application programs

By default, the language passed to a VTAM application program comes from the logon mode entry specified by the USS LOGMODE parameter. With multicultural support, the language passed to the application program can be specified by the user, independent of the LOGMODE parameter.

The user can enter a LOGON command to indicate which language is to be passed to the application program using:

- The language code parameter of the LANGTAB operand
- The LANG operand

Notes:

1. The language that is passed to the application program does not remain in effect across logons (as it does for USS messages). Each LOGON command must specify the LANG operand or the language code parameter of the LANGTAB operand to have a language passed to the application; otherwise, the default (MODEENT LANG operand value) is used.
2. If both operands indicate a language, the LANG operand has priority.
3. If the MMS is not currently active, the LANG operand is ignored.

For more information about the LANG and LANGTAB operands, see the [z/OS Communications Server: SNA Resource Definition Reference](#). For specific information about how the application program receives the specified language, see [z/OS Communications Server: SNA Programming](#).

Model name table

The model name table contains model names that can be passed to VTAM application programs in their LOGON exits. The purpose is to help VTAM application programs create dynamic definitions for their session-partner resources.

A model definition is normally used as the starting point for a dynamic resource definition. A typical application might maintain many different model definitions, one for each possible combination of resource characteristics. The model name passed by VTAM enables the application program to and accurately select the proper model definition to use.

Operands on an SLU resource definition can associate that SLU with the proper model name data. The MDLTAB operand specifies the model name table to be used, and the MDLENT operand specifies the proper entry within the table.

You create a model name table by specifying:

- A MDLTAB macro instruction
- One or more MDLENT macro instructions
- Optional MDLPLU macro instructions for each MDLENT

File these macros in the VTAM definition library. The name of the table is the name of the stored member or file. No assembly or link-edit is required to install the table. The table is built dynamically upon the first activation of any resource that has a defined association to the table.

IBM does not supply a default model name table.

An example of a model name table is shown in [Table 12 on page 202](#).

<i>Table 12. Example of model name table</i>		
Name	Operation	Operands
MDLTAB1	MDLTAB	
ENTRY1	MDLENT	MODEL=MDLNAME1
ENTRY2	MDLENT	
	MDLPLU	PLU=APPL2, MODEL=MDLNAME2
	MDLPLU	PLU=APPL5, MODEL=MDLNAME3
ENTRY3	MDLENT	MODEL=MDLNAME3
	MDLPLU	PLU=APPL5, MODEL=MDLNAME4
	MDLPLU	PLU=APPL7

Following are some examples of how the model name table in [Table 12 on page 202](#) works:

- Assume that the SLU has the MDLTAB=MDLTAB1 macro coded in its definition. The entry name specified by the SLU and the name of the PLU involved in the session determines which model name is sent to the PLU during session initiation.
- Assume the SLU has the MDLENT=ENTRY1 macro coded in its definition. This macro specifies a default model name of MDLNAME1 for the entry. Because no MDLPLU macros are specified, model name MDLNAME1 is sent for any PLU.
- Assume the SLU has the MDLENT=ENTRY2 macro coded in its definition. This macro specifies no default model name. The first MDLPLU macro specifies that model name MDLNAME2 is sent when the PLU is APPL2. The second MDLPLU macro specifies that model name MDLNAME3 is sent when the PLU is APPL5. For any other PLU, no model name is sent.
- Assume the SLU has the MDLENT=ENTRY3 macro coded in its definition. This macro specifies a default model name of MDLNAME3 for the entry. The first MDLPLU macro specifies that model name MDLNAME4 is sent when the PLU is APPL5. The second MDLPLU macro specifies that no model name is sent when the PLU is APPL7. For any other PLU, model name MDLNAME3 is sent.

Associated LU table

An associated LU table contains associated LU names that can be passed to VTAM application programs in their LOGON exits. The purpose is to help VTAM application programs create dynamic definitions for their session-partner resources.

The associated LU names provide supplementary information unavailable from model definitions. The LU names specify primary and alternate printers logically related to the SLU. (For information about model definitions, see [“Model name table” on page 201](#).)

Operands on an SLU resource definition can associate that SLU with the proper associated LU data. The ASLTAB operand specifies the associated LU table to be used, and the ASLENT operand specifies the proper entry within the table.

You create an associated LU table by specifying:

- An ASLTAB macro instruction
- One or more ASLENT macro instructions

- Optional ASLPLU macro instructions for each ASLENT

File these macros in the VTAM definition library. The name of the table is the name of the stored member or file. No assembly or link-edit is required to install the table. The table is dynamically built upon the first activation of any resource that has a defined association to the table.

IBM does not supply a default associated LU table.

Note: A printer must be in the same network as the terminal or must have a unique name that requires no translation across network boundaries. Violation of this restriction can cause failure of print-screen operations during cross-network sessions.

An example of an associated LU table is shown in [Table 13 on page 203](#).

Table 13. Example of associated LU table		
Name	Operation	Operands
ASLTAB1	ASLTAB	
ENTRY1	ASLENT	PRINTER1=PTRA,PRINTER2=PTRB
ENTRY2	ASLENT	
	ASLPLU	PLU=APPL2, PRINTER1=PTRC,PRINTER2=PTRD
	ASLPLU	PLU=APPL5,PRINTER1=PTRA
ENTRY3	ASLENT	PRINTER1=PTRA
	ASLPLU	PLU=APPL2, PRINTER1=PTRC,PRINTER2=PTRD
	ASLPLU	PLU=APPL7,PRINTER2=PTRB

Assume that the SLU has the ASLTAB=ASLTAB1 macro coded in its definition. The entry name specified by the SLU and the name of the PLU involved in the session determines which printer names are sent to the PLU during session initiation. For example:

- Assume the SLU has the ASLENT=ENTRY1 macro coded in its definition. This macro specifies a default primary printer name of PTRA and a default alternate printer name of PTRB for the entry. Because no ASLPLU macros are specified, primary printer name PTRA and alternate printer name PTRB are sent for any PLU.
- Assume the SLU has the ASLENT=ENTRY2 macro coded in its definition. This macro specifies no default printer names for the entry. The first ASLPLU macro specifies that primary printer name PTRC and alternate printer name PTRD are sent when the PLU is APPL2. The second ASLPLU macro specifies that primary printer name PTRA is sent and no alternate printer name is sent when the PLU is APPL5. For any other PLU, no printer names are sent.
- Assume the SLU has the ASLENT=ENTRY3 macro coded in its definition. This macro specifies a default primary printer name of PTRA for the entry. The first ASLPLU macro specifies that primary printer name PTRC and alternate printer name PTRD are sent when the PLU is APPL2. The second ASLPLU macro specifies that no primary printer name is sent and alternate printer name PTRB is sent when the PLU is APPL7. For any other PLU, primary printer name PTRA is sent and no alternate printer name is sent.

Selecting session parameters for the logon mode table

When a logical unit requests a session with an application program, it uses a symbolic logon mode name, either directly or by default, to suggest the session protocols. Session protocols are a set of rules that describe how the session is conducted. For example, one protocol might specify the application program uses chaining for SNA requests. Another might require the logical unit does not send end-bracket indicators to the application program. For a complete description of the session protocols available in VTAM, see [z/OS Communications Server: SNA Programming](#).

Session protocols are expressed as a string of characters called session parameters, usually specified in the logon mode table. A set of session parameters for a session is also called its logon mode. VTAM contains an IBM-supplied logon mode table, ISTINCLM. This table contains a set of generally accepted session parameters for a basic list of IBM device types, but it might not completely meet your needs. You can create supplementary logon mode tables and associate them with device type logical units or application programs using the MODETAB operands on the definition statements defining them.

To create the logon mode table:

- Modify the IBM-supplied table.
- Create supplementary tables and associate them with device-type LUs or application programs using the MODETAB operands in the definition statements defining them. You can use the DLOGMOD operand on the appropriate definition statements to indicate the specific logon mode entry in the table that contains the session parameters to be used by default. The LOGMODE operand can also be specified on VARY LOGON commands to specify a logon mode table entry.
- Associate a supplementary table with dynamic CDRSCs using the DYNMODTB start option. You can also use the DYNDLGMD start option to indicate the default logon mode entry for dynamic CDRSCs. The value for DYNMODTB can be changed using the DYNMODTB operand on the MODIFY VTAMOPTS command, or by using the MODIFY TABLE,OPTION=LOAD command to replace the existing table that is being used by dynamic CDRSCs. You can change the default logon mode entry for dynamic CDRSCs using the DYNDLGMD operand on the MODIFY VTAMOPTS command. You can also override this value for a specific dynamic CDRSC using the MODIFY DEFAULTS or MODIFY RESOURCE commands.

When VTAM builds a dynamic CDRSC, it uses the DYNMODTB value (if specified; otherwise ISTINCLM) to associate a logon mode table to the dynamic CDRSC. It also uses the DYNLOGMD value, if specified, to assign a default logon mode table entry to the dynamic CDRSC. Changing the value of these start options does not change the value for dynamic CDRSCs that have already been built.

Logon mode names are always resolved using the logon mode table associated with the secondary logical unit (SLU). For more information about which node performs the mode resolution, see [“Resolving logon mode names to subarea and APPN Classes of Service”](#) on page 417.

Logon mode tables are created or modified using MODETAB, MODEENT, and MODEEND macro instructions. The [z/OS Communications Server: SNA Resource Definition Reference](#) describes the contents of the IBM-supplied logon mode table. It also describes how TSO can query for the logon mode being used by a terminal in place of supplying the LOGMODE during the logon. [Figure 55 on page 205](#) shows the sequence in which MODETAB, MODEENT, and MODEEND macros are coded to define a logon mode table.

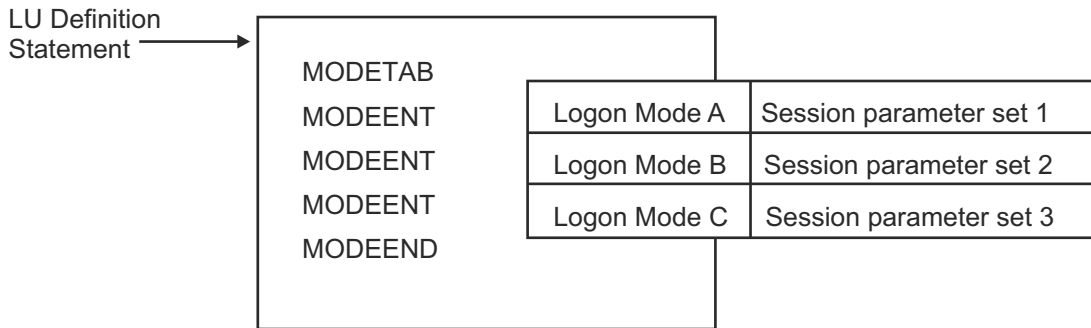


Figure 55. Macroinstructions for logon mode table

The various ways of specifying session parameters are illustrated in [Table 14](#) on page 205.

Table 14. How session parameters are identified			
Request source	Request form	How session parameters are identified	
		Logon mode table entry	Logon mode table
Dependent Logical Unit	<ul style="list-style-type: none"> Initiated Self request Interpret table character string representing logon Character-coded logon 	<ul style="list-style-type: none"> Field in Initiate Self request containing mode name used to search logon mode table Parameter in a character-coded logon that translates into LOGMODE (name) DLOGMOD operand on LU definition statement Default entry ISTDOSDF, if the ISTDOSDF start option allows its use for dependent LUs First entry in logon mode table named in MODETAB operand on LU definition statement First entry in IBM-supplied logon mode table 	<ul style="list-style-type: none"> MODETAB operand on LU definition statement IBM-supplied logon mode table

Table 14. How session parameters are identified (continued)

Request source	Request form	How session parameters are identified	
		Logon mode table entry	Logon mode table
Independent Logical Unit	<ul style="list-style-type: none"> • BIND • Cross-domain initiate 	<ul style="list-style-type: none"> • Field in BIND or cross-domain initiate request containing mode name used to search logon mode table • DLOGMOD operand on CDRSC definition statement, or DYNDLGMD start option value (for dynamic CDRSCs) • Default entry ISTCOSDF, if the ISTCOSDF start option allows its use for independent LUs • First entry in logon mode table named in MODETAB operand on CDRSC definition statement or DYNMODTB start option value (for dynamic CDRSCs) • First entry in IBM-supplied logon mode table 	<ul style="list-style-type: none"> • MODETAB operand on CDRSC definition statement or DYNMODTB start option value (for dynamic CDRSCs) • IBM-supplied logon mode table
Application Program	Application program issuing VTAM macros, such as SIMLOGON or REQSESS	<ul style="list-style-type: none"> • Name of logon mode table entry specified in the NIB • DLOGMOD operand on LU definition statement • Default entry ISTCOSDF, if the ISTCOSDF start option allows its use for application programs • First entry in logon mode table named in MODETAB operand on LU definition statement • First entry in IBM-supplied logon mode table 	

Table 14. How session parameters are identified (continued)			
Request source	Request form	How session parameters are identified	
		Logon mode table entry	Logon mode table
Automatic logon	Initiate Other request	<ul style="list-style-type: none"> • DLOGMOD operand on LU or CDRSC definition statement or DYNDLGMD start option value for dynamic CDRSCs • Default entry ISTCOSDF, if the ISTCOSDF start option allows its use • First entry in logon mode table named in MODETAB operand on LU or CDRSC definition statement, or DYNMODTB start option value for dynamic CDRSC • First entry in IBM-supplied logon mode table 	
VTAM operator	VARY LOGON command	<ul style="list-style-type: none"> • LOGMODE operand in VARY LOGON command • LOGMODE operand in previous VARY LOGON command • DLOGMOD operand on LU or CDRSC definition statement, or DYNDLGMD start option value for dynamic CDRSCs • Default entry ISTCOSDF, if the ISTCOSDF start option allows its use • First entry in logon mode table named in MODETAB operand on LU or CDRSC definition statement, or DYNDLGMD start option value for dynamic CDRSCs • First entry in IBM-supplied logon mode table 	

Regardless of the source of the logon or what session parameters are associated with the logon, the application program decides which session parameters are used for the session. It can decide to use the

session parameters associated with the pending logon, or it can choose a different set of parameters. When the application program issues a macro to initiate a session with a logical unit, the application program indicates the logon mode name or session parameters to be used during the session. The logon mode table from which the session parameters are selected is always the one associated with the SLU, or, if a table is not associated with the logical unit, it is the IBM-supplied default logon mode table in the VTAM that owns the SLU.

If the specified logon mode entry does not exist in the logon mode table associated with the secondary logical unit, the default logon mode (ISTCOSDF) is used if it is found in the table and its use is allowed as determined by the ISTCOSDF start option.

You can modify or replace the IBM-supplied logon mode table, provided that the modified or replacement table has the same name as the IBM-supplied table and that the IBM-supplied table is deleted. However, because the IBM-supplied table might be needed for problem determination, you should create supplementary tables instead of deleting the IBM-supplied table.

As part of the session-establishment procedures, an application program acting as the PLU can:

- Supply a logon mode name associated with the SLU. The session parameters associated with the logon mode name in the domain of the SLU are used.
- Directly supply the session parameters to be used.
- Use the session parameters associated with the pending logon.

[z/OS Communications Server: SNA Programming](#) describes in more detail how the application program handles session parameters.

[Figure 56 on page 209](#) shows a summary of the algorithm VTAM uses to obtain session parameters.




Figure 56. How session parameters are obtained from a logon mode table

Automatic logons

You can enable your LUs to automatically log on to a particular application program. This is most useful when a designated terminal always accesses one application program. To enable automatic logons, you can either code the LOGAPPL operand on the LU definition statement, or you can use the VARY LOGON or VARY ACT,LOGON commands.

Automatic logon is useful for the following situations:

- Unattended terminals (terminals with no operator to issue a logon)
- Establishing a controlling application program to:
 - Manage allocation of terminals
 - Automatically reestablish sessions with terminals after failure and restart of a host or NCP

Coding for automatic logon

By using the LOGAPPL operand on the definition statement for an LU, the LU can automatically initiate a session with a specific application program whenever the LU is activated. The application program can either be the program with which the LU ultimately initiates a session or an intermediary application program to which the LU is temporarily assigned. For application programs in another network, the network-qualified name can be specified on the LOGAPPL operand.

The application program to which an LU is automatically logged on is called its controlling application program. If the LU is passed to another application program, VTAM automatically initiates a session with the controlling application program when that session is terminated. If the application program with which the LU is in session terminates that session and there is no controlling application program, the LU is available for session initiation.

[Figure 57 on page 211](#) shows a sample automatic logon session:




Figure 57. Automatic login to A50ACCTS application program

See the following sample coding for the A50ACCTS application program in [Figure 57 on page 211](#):

```
APPLNODE VBUILD TYPE=APPL
A50ACCTS APPL  ATNLOSS=ALL,    ENABLE ATTN EXIT FOR ALL LU 6.2 SESSIONS
                APPC=YES,      ENABLE LU 6.2 SUPPORT
                PARSESS=YES,    ENABLE PARALLEL SESSION FOR LU 6.2 SUPPORT
                SECACPT=ALREADYV, ACCEPT ALREADY VERIFIED LU 6.2 CONVERSATIONS
                ACBNAME=ACCOUNTS NETWORK LU NAME=A50ACCTS
                VTAM APPLICATION APPLID=ACCOUNTS
                :
```

See the following sample coding for the 3277 display station in [Figure 57 on page 211](#):

```
LC3270NS LBUILD          LOCAL NON-SNA ATTACHMENT
LOC3277  LOCAL  CUADDR=00A, CHANNEL UNIT ADDRESS
          :      TERM=3277,  TERMINAL TYPE
          :      DLOGMOD=S3270, DEFAULT LOGON MODE TABLE ENTRY
          :      LOGAPPL=A50ACCTS AUTOMATIC LOGON APPLICATION
```

Operator commands for automatic logon

You can establish an automatic logon session using the VARY LOGON command. After the automatic logon is established, it remains until one of the following conditions occurs:

- The major node of one of the session partners is deactivated.
- The primary logical unit is changed using the VARY LOGON or VARY ACT,LOGON command. If the primary logical unit is a USERVAR name, the session setup requests can be rerouted to a different primary logical unit by using the MODIFY USERVAR command.
- The automatic logon is deleted using the VARY NOLOGON command.

Note:

1. You can view the automatic logon relationship of an LU with the DISPLAY ID=*sluname* command.
2. You can view automatic logon relationships by controlling application name using the DISPLAY AUTOLOG,ID=*controlling_application_name* command. This display shows only LUs that are session capable and not currently in session with their controlling application. This display also shows the events that can recover the controlling sessions.
3. The VARY AUTOLOG command can be used to attempt recovery of the controlling sessions.
4. See [z/OS Communications Server: SNA Operation](#) for more details on using the DISPLAY AUTOLOG and VARY AUTOLOG operator commands.

Reallocation of autologon sessions

For dependent LUs and TN3270 applications using DEFAPPL, the VTAM that owns the secondary logical unit controls the automatic logon sessions. For independent LUs, the VTAM providing boundary function services for the secondary logical unit controls the automatic logon sessions. VTAM also provides autologon support over a CP-CP session for secondary LUs whose network node server does not have autologon capabilities.

By default, adjacent CDRM activation redrives any autologon sessions that are not established between resources of the two VTAMs. Adjacent CP activation also redrives any such autologon sessions, if the adjacent CP supports automatic logon.

To specify which, if any, adjacent node activations are to result in trying pending autologon requests again, use the AUTORTRY start option. In addition to the default for AUTORTRY described above, you can use AUTORTRY to limit session tries again to only after adjacent CDRM activation (AUTORTRY=CDRM), to suppress all retries for adjacent node activations (AUTORTRY=NONE), or to specify that all adjacent node activations are to result in session retry (AUTORTRY=ALL). For AUTORTRY values other than NONE, CDRM deactivation causes redrives of any autologon sessions that relied on the lost CDRM. CP deactivation does not cause redrives.

Use the AUTOTI start option to control how often pending autologon requests are retried. Time values are specified in seconds, minutes, hours, or days. If AUTOTI is nonzero, redrives occur at timer expiration for all autologon sessions that are not established, regardless of the value of AUTORTRY. Use AUTOTI to ensure session setup in the following situations:

- Multihop CP-CP network
- Network node server that has not yet received the controlling primary LU registration
- CLSDST PASS that was sent to an unavailable application
- USERVAR value that changed while a node is inactive, and the node is then reactivated

- Controlling PLU that was a clone application that had not yet opened its ACB
- Controlling PLU was a clone application that has moved to a new host (for example, CLOSE ACB on HOST A and then OPEN ACB on Host B)

Notes:

1. When using search reduction support (SRCHRED=YES, and SRTIMER, SRCOUNT, or both are nonzero) in conjunction with a nonzero value for AUTOTI, remember that the two functions are designed to produce opposite effects. Search reduction is used to limit searches for resources that have been found to be unreachable. AUTOTI is used to intermittently search the network at a specified interval for a previously unreachable controlling PLU. After the first search for a controlling PLU fails, attempts are made to redrive unestablished autologon sessions whenever AUTOTI expires. However, if AUTOTI is less than SRTIMER, these requests for the PLU are limited by the search reduction entry until it also expires. If AUTOTI=x and SRCOUNT=z, such requests for the PLU are limited for (x*z) seconds. For information about search reduction, see [“Improving VTAM performance using start options”](#) on page 24.
2. The VARY AUTOLOG operator command can be used to manually reallocate autologon sessions.

Sample automatic logon reallocation

In [Figure 58](#) on page 213, assume that secondary logical unit LU1A is in session with an application program APPL1, and that LU1A has LOGAPPL=A50ACCTS coded. When the session with APPL1 ends, VTAM attempts to establish a session between LU1A and A50ACCTS.

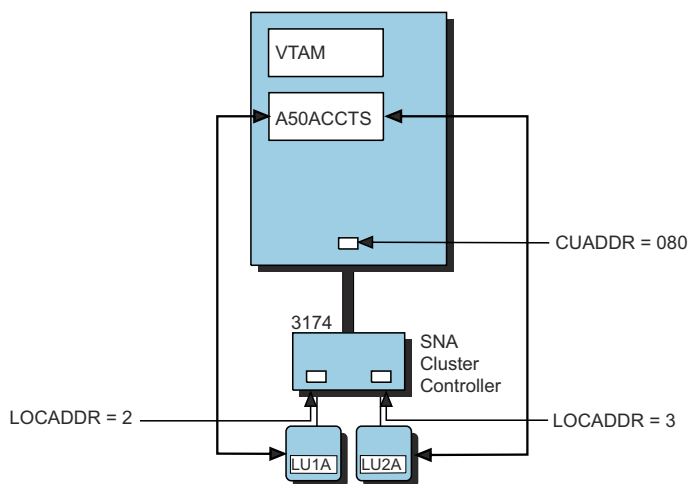


Figure 58. Automatic logon reallocation

If A50ACCTS is inactive, the automatic logon session fails. If the A50ACCTS major node is active to VTAM but it has not opened its ACB yet, the automatic logon session fails, but is initiated when A50ACCTS opens its ACB.

If you use the VARY INACT command to deactivate A50ACCTS, you can only regain the automatic logon sessions with LU1A and LU2A by issuing a VARY ACT,LOGON command for A50ACCTS or reactivating A50ACCTS and then reactivating LU1A and LU2A.

Conditions for reallocation

A session is reallocated as a result of one of the following conditions:

- If a secondary logical unit goes through a state change from disabled or inhibited to enabled, the owning VTAM attempts to set up a session with the controlling application program for the secondary logical unit. The session state is determined by the device and is indicated to VTAM by an ACTLU response from the logical unit or a NOTIFY request from the device.

- If a session is terminated, VTAM checks to see whether an automatic logon relationship exists and, if so, reallocates that session if the session being ended is not with the controlling application program.
- When an active session between a secondary logical unit and its controlling application program terminates, the VTAM that owns the secondary logical unit decides whether to reallocate the controlling session based on the type of termination request or UNBIND that is received. The controlling session is not reallocated if the termination request or UNBIND type indicates normal termination. If the session was terminated using the VARY TERM command with either TYPE=UNCOND or TYPE=FORCE, the controlling session is reallocated if a session path exists at that time between the two logical units. To avoid redriving the session, issue VARY TERM TYPE=COND instead. For further information, see [z/OS Communications Server: SNA Operation](#).

How primary logical unit status affects reallocation

VTAM attempts to start a session with the controlling application program (the primary logical unit). Following is an explanation of how VTAM reacts to the primary logical unit status:

- If the owning host of the secondary logical unit cannot find the primary logical unit (the major node is not active) or is inactive, the automatic logon session fails. VTAM does not automatically reallocate the failing session to the primary logical unit; instead, the owning host of the secondary logical unit creates a pending autologon request. VTAM retries pending autologon requests periodically based on the settings of the AUTOTI and AUTORTRY start options.
- If the primary logical unit is in CONCT status (active to VTAM but OPEN ACB has not been issued), the automatic logon fails with sense code 088A0003, and VTAM tries to set up the session again when the primary logical unit issues OPEN ACB and SETLOGON OPTCD=START macroinstructions.
- If the primary logical unit is deactivated through VTAM by an operator command and becomes inactive, the secondary logical unit must reattempt the automatic logon session using a secondary logical unit state change, or the operator can issue a VARY ACT command for the secondary logical unit after the primary logical unit is at least in CONCT status.

Note: If an SLU is in session with its controlling application at the time of the PLU deactivation, VTAM retries the automatic logon session. If VTAM retries to establish a session and fails, a pending autologon request is created in the owning host of the SLU. VTAM retries pending autologon requests periodically based on the settings of the AUTOTI and AUTORTRY start options.

Determining automatic logon relationships

To determine an automatic logon relationship, use the DISPLAY ID=(*SLUname*),E command. Use this command to determine whether a logical unit or TN3270 application has a controlling primary logical unit. Message IST1131I contains the name of the primary logical unit. Message IST635I contains the name of the session partner, the status of the session, and other session information. The status of the session will have the last two characters overlaid with /C if this session is because of an automatic logon relationship.

The DISPLAY AUTOLOG command displays controlling applications for which there are one or more LUs in this host that are session capable but that are not currently in session with their controlling application. The command can also be used to list SLUs that are in this state. The PLU name is displayed in messages that include the events that automatically attempt to re-establish sessions between the controlling application and the SLUs that have pending automatic logon requests.

Sample displays received during reallocation

If the primary logical unit is in CONCT status when the automatic logon session setup fails, you will receive the following messages:

```
IST663I INIT OTHER OR CDINIT REQUEST FAILED, SENSE=088A0003
IST664I REAL OLU= <SLU-NAME> REAL DLU= <PLU-NAME>
IST889I SID= <SESSION ID>
IST890I AUTOLOGON SESSION SETUP FAILED
IST896I AUTOLOGON WILL BE RETRIED WHEN CONTROLLING PLU IS AVAILABLE
IST314I END
```

Note: Besides normal message suppression, the number of IST663I and IST664I messages for the 088A0003 sense codes can be suppressed using the PLUALMSG and SLUALMSG start options. For more information, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

If the primary logical unit was in CONCT status when the automatic logon failed and then issues an OPEN ACB, you will receive the following messages:

```
IST899I RETRY OF AUTOLOGON(S) TO <PLU-NAME> IN PROGRESS
```

If the primary logical unit is not active when the automatic logon fails, you will receive the following messages:

```
IST663I INIT OTHER OR CDINIT REQUEST FAILED, SENSE=08570002
IST664I REAL OLU= <SLU-NAME> REAL DLU= <PLU-NAME>
IST889I SID=.....
IST890I AUTOLOGON SESSION SETUP FAILED
IST1138I REQUIRED RESOURCE <PLU-NAME> NOT ACTIVE
IST314I END
```

Note: The sense code in message IST663I, and the text of message IST1138I might be different.

If you had automatic logon sessions fail on a primary logical unit that was in CONCT status and then you deactivate it before it opens its ACB, you will receive the following messages:

```
IST097I VARY ACCEPTED
IST899I RETRY OF AUTOLOGON(S) TO <PLU-NAME> WILL NOT OCCUR
IST1133I <PLU-NAME> NODE NOW INACTIVE
```

If the VARY AUTOLOG command is used to attempt reallocation of automatic logon sessions, you will receive the following messages:

```
IST097I VARY AUTOLOG ACCEPTED
IST2093I AUTOLOGON SEARCH INITIATED FOR <NUMBER> APPLICATIONS
IST314I END
```

At this point, searches are generated to find and determine the status of one or more controlling applications. After the search for each controlling application completes, the following messages are issued if sessions cannot be established with the controlling application.

- If the search is unable to locate the controlling application, the following messages are issued:

```
IST2096I AUTOLOGON SEARCH COMPLETED FOR APPLICATION <PLU-NAME>
IST2097I STATUS - CANNOT BE LOCATED
IST314I END
```

- If the search locates the controlling application and finds the status of the application is not session capable, the following messages are issued:

```
IST2096I AUTOLOGON SEARCH COMPLETED FOR APPLICATION <PLU-NAME>
IST2098I STATUS - UNABLE TO ACCEPT LOGONS
IST314I END
```

- If the search finds the controlling application capable of accepting sessions, the following messages are issued:

```
IST2096I AUTOLOGON SEARCH COMPLETED FOR APPLICATION <PLU-NAME>
IST2099I STATUS - AUTOLOGON SESSIONS INITIATED FOR <NUMBER> LUS
IST314I END
```

Session management exits

The session management exit routine can be used to:

- Check or restrict the use of a logical unit (session authorization)
- Gather accounting information (session accounting)
- Select adjacent SSCP for LU-LU session and resource status requests

- Process a session takeover by the alternate XRF application program

Code a session management exit routine if you want to use these functions. There is no default session management exit routine. If you do not code a session management exit routine, it is assumed that all session requests are authorized and no accounting data is collected by VTAM.

VTAM calls the session management exit routine whenever a session between logical units is established, ended, taken over, or recovered. The session management exit routine is invoked by every VTAM in the session setup path. It is also called for session accounting when an LU-LU session ends.

Note: Improper use of the session management exit routine can impact VTAM operation. For example, you should consider the following possibilities:

- A program error can affect VTAM, the operating system, and other programs. Although abnormal end recovery is provided for errors within the exit, the exit code can accidentally corrupt VTAM storage.
- Lengthy processing time could degrade VTAM performance. While these routines are running, VTAM cannot process any new VTAM operator requests to initiate or end sessions. To reduce the effect on VTAM performance, any analysis of the information collected should be done later by another program.

In a multiple domain environment, the session management exit routine can be invoked in the domains in which each of the session partners resides and can be used to select adjacent SSCPs for LU-LU session and resource status requests.

The adjacent SSCP selection function in the session management exit routine can be used to control the SSCP to which a session request is sent. The exit allows you to shorten or reorder the list of adjacent SSCPs, or to not route at all. You can use this function with the dynamic adjacent SSCP table facility so that you can avoid defining adjacent SSCP tables while maintaining control of session request routing.

VTAM supports session authorization exit routines, which can be used to control session initiation. These routines can accept or reject the session based on various information that is passed to them. VTAM also supports session accounting exit routines that can be used to perform accounting functions for established sessions.

The session authorization and session accounting exit routines provide some of the same functions as the session management exit routine. These exits are invoked before the corresponding session management function. However, these exit routines do not provide all of the functions that are available in the session management exit routine, and they are not recommended.

For further information about the session management, session authorization, and session accounting exit routines, including a detailed description of the vector lists and the information provided to the routines, see [z/OS Communications Server: SNA Customization](#).

Session authorization

When SSCP routing is in process for cross-domain sessions, full information about the destination logical unit might not be available when the session management exit routine is first invoked. To enable proper authorization to be done for cross-domain sessions, two separate calls for authorization can be made for these sessions. The first call is made early in the session setup process. The session management exit routine can accept or reject the session at this time, or it can defer its decision until later. If the decision is deferred, the routine is called again for a secondary authorization after full information is known about the destination logical unit.

You might code the routine to contain a table of valid sessions against which the session-establishment request can be compared. For example, an application program can be designed to establish a session with any logical unit using the OPNDST macro instructions with the ACCEPT option in its LOGON exit routine. The authorization exit routine can compare the identity of any logical unit that attempts to establish a session with the application program to entries in such a table to determine whether authorization can be granted for that logical unit.

Both primary and secondary authorization exits are also driven during LU takeover processing if the sessions being taken over were established using extended BIND protocols.

Session accounting

With the capability of peripheral LUs serving as primary LUs (PLUs), accounting application programs relying on the primary LU to be located only in the host might not be sufficient.

At LU-LU session setup time, a session awareness PIU is sent to the NetView Performance Monitor (NPM) as unsolicited information. Byte and PIU counts are updated every time there is traffic on the session. When the counts surpass the user-specified threshold, the counters are sent to NPM. At the end of the session, session awareness data and the last set of counters for the session are sent to NPM. The user-specified options, and thresholds, can be changed using NPM commands. In addition, transport of accounting data to NPM can be suspended and resumed using NPM commands.

If selecting session accounting, the user can specify the following during NCP generation:

- Whether session accounting is required for PLUs, SLUs, or all LUs
- Whether accounting collection should begin immediately or be deferred
- The session accounting byte and PIU thresholds
- Whether backup NPM sessions are defined
- Whether gateway session accounting is to be used for cross-network sessions
- The number of half-sessions for which session accounting will be done
- The number of session accounting extension blocks to define
- The PIU distribution ranges (if any) to be used

Especially important to a multiple-network environment is the ability of the user to specify the following during NCP generation:

- Whether gateway session accounting is to be used for cross-network sessions
- The number of half-sessions for which session accounting is done

Session-level pacing

Session pacing can be used to influence performance between session partners. Using high-pacing window sizes or no-session pacing for interactive sessions and using lower-session pacing window sizes for batch sessions is a widely used approach to begin tuning an SNA network. For example, a VTAM application driving a low-speed printer needs to use small session pacing window sizes to keep one session partner from overwhelming the other with data. Using higher (or no) pacing window sizes for interactive sessions than for batch sessions favors the interactive users in the network containing both batch and interactive traffic. For batch, if the session partners and network can support it, higher-pacing window sizes can be used to improve throughput.

Session-level pacing involves each session-layer component in a session path. Session pacing occurs in session stages. Each stage is independently paced, allowing different pacing window sizes to be used for each stage.

There are two types of session pacing:

Fixed

With fixed session-level pacing, the pacing window size is constant. The amount of data that can be transmitted or received between session partners is maintained as a constant window size.

Adaptive

With adaptive session-level pacing, the window size can vary depending on the resources available at each session endpoint. After each window of data is transmitted, the receiver informs the transmitter how large the next window can be.

Adaptive pacing will be used whenever possible. A session path between two LUs may have a mixture of pacing types along the route.

See the *SNA Technical Overview* for additional information about session-level pacing concepts.

Fixed session-level pacing

The pacing window size between two endpoints of a pacing stage is constant when fixed session-level pacing is being used. The window size is passed in the BIND request during session activation. The window size in the BIND response is the actual window size used. The number of pacing stages can be the same for both directions (primary-secondary and secondary-primary) in a session, or it can be different. In [Figure 59 on page 218](#), APPL1 is in session with a type 2.0 node. Two-stage pacing occurs from APPL1 (primary logical unit) to the type 2.0 node (secondary logical unit), and one-stage pacing occurs from the secondary logical unit to APPL1. In this case, the NCP provides the boundary function for the LU.




Figure 59. Fixed-session pacing (one- and two-stage)

For fixed session pacing, the pacing response is sent in either of two ways:

- In the response header of a message unit
- In a stand-alone message called an isolated pacing response (IPR)

Adaptive session-level pacing

Adaptive session pacing allows the pacing window size to dynamically:

- Increase to accommodate increased traffic
- Decrease when congestion is in the receiving node

The pacing window sizes are adapted to buffer availability and demand on a session-by-session basis. The session stage endpoints exchange explicit pacing windows that can vary in size during the course of a session. The mechanism thus prevents any single session from monopolizing or exhausting the buffer resources of a session stage endpoint.

Adaptive pacing occurs between each pacing stage endpoint and can vary for each direction. In [Figure 60 on page 219](#), APPL1 is in session with a type 2.1 node. Two-stage pacing occurs for data flowing in both directions between the session partners.

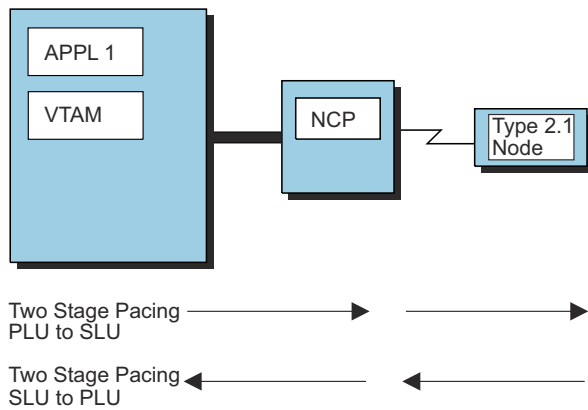


Figure 60. Adaptive session pacing

If adaptive pacing is performed, the value determined during session setup for the pacing window size represents the minimum pacing window size to be used during normal operation.

Note: Before VTAM Version 4 Release 4, VTAM did not dynamically vary the session partner send window size. VTAM always returns the same window size that was set during session establishment. However, VTAM honors the pacing windows sent by its pacing stage partners for adaptively paced sessions. The next window size used for data transmission is whatever value was received from the session partner in the pacing response.

An isolated pacing message (IPM) indicates whether the pacing window size should be changed. Pacing window sizes are changed by IPMs. The IPM, like the IPR used in fixed session pacing, is used to authorize the sending of the next window of messages units, but it also specifies whether the window size is to be changed or remain the same.

The following scenario describes how adaptive pacing works.

1. The transmitter requests a pacing response from the receiver when sending the first PIU in the current send window.
2. The receiver returns a pacing response that contains the next window size it considers appropriate for the next transmission. If the transmitter sends the entire window of PIUs before the pacing window response is received, the transmitter must wait until the pacing response is received.
3. The transmitter can request a larger window size by setting an indicator in the next transmission.
4. The receiver can ignore the transmitter request for a larger window or honor it when generating the pacing response.

The transmitter is obligated not to exceed the window size specified by the receiver.

Note: You must have NCP Version 5 Release 2, or higher, to support two-stage and adaptive-session pacing involving a 3745 communication controller, or Version 4 Release 3, or higher, for a 3725 communication controller.

Setting initial pacing values

The type of session-level pacing and pacing window sizes are determined by the BIND request when a session is established. When an extended BIND is used, adaptive session-level pacing is requested. If adaptive-session pacing is not supported or required by the session partner, the BIND response will be sent indicating fixed-session pacing. (XBSI is the extended BIND indicator in the session initiation request. XBSI=ON means that extended BIND is used to establish the session.) The actual values used are those in the BIND response.

The pacing values are carried in the BIND request flowing from the primary LU to the boundary function (if one exists for this session), and then to the secondary LU. In this way, all the participants in the session are informed of the pacing values to be used.

With this in mind, the possible pacing windows for a session are:

PS

Primary send

SR

Secondary receive

SS

Secondary send

PR

Primary receive

The direction of the data being sent determines which pacing window values are used.

- When data is sent from the primary LU to the secondary LU, the PS and SR values are used.




Pacing Response

Message unit 1

Message unit 2

Message unit 3
Message unit 4

- The primary send count controls the flow of requests from the primary logical unit to the boundary function,
- The secondary receive count controls the flow from the boundary function to the secondary logical unit.
- When data is sent from the secondary LU to the primary LU, the SS and PR values are used.



- The secondary send count controls the flow of requests from the secondary logical unit to the boundary function.
- The primary receive count controls the flow from the boundary function to the primary logical unit.

The values for the four pacing windows are derived from:

- The following operands on the MODEENT macro instruction for the logon mode entry:

- PSNDPAC for primary send
- SRCVPAC for secondary receive
- SSNDPAC for secondary send
- The following operands on the LU or CDRSC definition statement:
 - PACING
 - VPACING
- The following operands on the APPL definition statement:
 - AUTH=NVPACE|VPACE
 - VPACING

Note: An application program acting as a primary end of a session can obtain information about pacing counts using the INQUIRE macro instruction or the LOGON exit routine. For more information, see [z/OS Communications Server: SNA Programming](#).

The sources used to determine pacing-window values are dependent on the configuration and the secondary LU (SLU). [Figure 61 on page 222](#) and [Figure 62 on page 222](#) describe the environment, and [Table 15 on page 221](#) identifies the sources used in determining the initial pacing window value when the primary LU (PLU) is an application program. For example, if the SLU is an application program, the pacing sources, regardless of whether an extended BIND is sent, are described in [Table 15 on page 221](#) under:

- Letter I for primary to secondary flow
- Letter D for secondary to primary flow

Notes:

1. For fixed-session pacing, the pacing window values derived are the actual values used.
2. For adaptive-session pacing, the pacing window values derived represent the minimum pacing window size to be used.

<i>Table 15. Correspondence of methods to letters</i>	
Letters	Methods used by VTAM
A	If SRCVPAC is not 0, use it else use SLU VPACING (default is 7)
B	If SRCVPAC is not 0, use it else use SLU PACING (default is 1)
C	Use SRCVPAC
D	If SSNDPAC is 0, use it else use PLU VPACING (default 7)
E	Use SSNDPAC
F	If PLU AUTH=NVPACE, use 0 else if PSNDPAC is not 0, use it else use SLU VPACING (default is 2)
G	If PLU AUTH=NVPACE, use X'5F' else if PSNDPAC is not 0, use it else use SLU VPACING (default 1)

Table 15. Correspondence of methods to letters (continued)

Letters	Methods used by VTAM
H	Use PSNDPAC
I	If PLU AUTH=NVPACE, use 0 else if SRCVPAC is not 0, use it else use SLU VPACING
J	Use PLU VPACING (default 7)

	If SNA SLU Is:											
	An Application				An NCP LU				A Same-Domain Local LU			
And the Pacing Direction Is:	P→S	S→P	P→S	S→P	P→S	S→P	P→S	S→P	P→S	S→P	P→S	S→P
And the Windows Are:	PS	SR	SS	PR	PS	SR	SS	PR	PS	SR	SS	PR
Windows Are Set Using Method:	I	D	F	B	D	A	D	G	B	E	J	

Figure 61. Pacing windows for SNA LUs

	If Non-SNA SLU Is:							
	A Same-Domain LU or an NCP Terminal				A Cross-Domain 370 CA Terminal			
And the Pacing Direction Is:	P→S	S→P	P→S	S→P	P→S	S→P	P→S	S→P
And the Windows Are:	PS=SR	SS=PR	PS	SR	SS=PR	PS=SR	SS=PR	
Windows Are Set Using Method:	C	E	H	C	E	C	E	

Figure 62. Pacing windows for non-SNA LUs

Overriding fixed pacing counts

The pacing counts specified on the LU definition statement apply to all sessions in which the LU participates. Values stated in the APPL definition statement override the values in the logon mode table entry and the LU definition statement. When you specify AUTH=NVPACE, the value for the primary send that goes out in the BIND is 0. Therefore, there is no pacing between the application program and the boundary function.

AUTH=VPACE prevents overloading buffers in the communication controller with outbound messages from VTAM. You can specify NVPACE when the application program sends a limited amount of data and then waits for a user response.

The primary application program can override the primary receive value for a specific session by specifying a nonzero value in the primary receive field of the session parameters it specifies at OPNDST. If 0 is specified, the primary receive value present in the session parameters in the pending logon is used. Only values less than or equal to the value in the pending logon should be specified by the application program.

Note: NVPACE can also be specified when an application program is sending chains of RUs to a secondary logical unit and the application program cannot send more BIUs than can be held in the queue. The number of elements in each chain must be no larger than the results of the following formula, where n is the smallest value in the VPACING operands of the definition statements for the LUs:

$$\text{number of elements} \leq 2n - 1$$

If the application program is in session with LUs whose VPACING operands are different from each other, the smallest VPACING values should be used in the formula.

See [“Setting initial pacing values” on page 219](#) for information about how VTAM determines the initial pacing-window size.

Independent logical units and adaptive-session pacing

Although some type 2.1 peripheral nodes do not support adaptive-session pacing, an independent logical unit in that node can still act as the primary LU and therefore sends the BIND. The primary pacing values in the BIND are established by the independent logical unit (primary). However, the value of the primary send count might be so large that it can impact the subarea node to which it is attached. If a type 2.1 peripheral node that supports adaptive session pacing is attached directly to VTAM, VTAM honors the pacing-window-size value included in the pacing response from the peripheral node but continues to use the appropriate receive count for its fixed-window-size value.

The primary send count specified in the logon mode entry of the primary logical unit is used by VTAM to control the rate at which a peripheral node acting as a primary logical unit sends data to an NCP. This value is then used by NCP or VTAM for a directly attached type 2.1 peripheral node to control the primary send pacing. If a negotiable BIND is used, the smaller value (other than 0) is used. If the BIND is nonnegotiable and the primary send value in the original BIND (other than 0) is larger than the value VTAM obtained from the logon mode entry or PACING specification, VTAM or NCP may reject the session setup request. For a directly attached type 2.1 peripheral node, VTAM performs the same type of operation.

Therefore, you must specify proper pacing values in VTAM for any type 2.1 independent logical unit. If the type 2.1 peripheral node supports adaptive-session pacing, this procedure is not necessary.

Application program pacing

This section describes how to use the APPL definition statement to specify pacing for application programs.

Inbound session pacing

Application programs can act as either primary or secondary LUs. If the application program is the primary logical unit, you define pacing with two values: a nonzero value in SSNDPAC on the secondary LU-mode table entry and the pacing value on the VPACING operand of the APPL definition statement. In this way, you can pace normal-flow requests from another logical unit.

If the application program is a secondary logical unit, you define pacing with either value: a nonzero value in PSNDPAC on the secondary LU-mode table entry, or the pacing value on the VPACING operand of the secondary logical unit APPL definition statement. If you code both values, the nonzero value in PSNDPAC overrides the VPACING value.

Outbound session pacing

When the application program is acting as the primary logical unit, you can specify AUTH=NVPACE to nullify pacing from that application program. VTAM does not wait for a pacing response before sending data to the NCP on behalf of this session. VPACE, the default, is normally used to prevent overloading buffers in the communication controller with outbound messages from VTAM. However, you can specify NVPACE when either of the following situations occur:

- The application program sends only single-element request units (RUs) to any one logical unit and, after sending each RU, waits for a response before sending the next RU.
- The application program sends chains of RUs containing a limited number of elements and does either of the following:
 - After sending one chain, the application waits for a response before it starts sending the next chain.
 - The application sends the change direction indicator in the last element actions of the chain.

Local flow control is a part of network flow control and can affect application program design. For primary or secondary application programs, macros that cause data to be sent on a session (SEND, SESSIONC, OPNDST, or CLSDST) are not posted complete when a virtual route (VR) is blocked because of congestion. The application program is posted when the VR is again opened. The data to be sent is left in application program buffers.

An application program that issues a synchronous SEND request does not receive control until the request is processed, even if the route is not blocked. For an asynchronous request, the program receives control as soon as the request is accepted by VTAM. However, further processing of the SEND RPLs for this session is suspended as long as the VR is blocked.

Sample configurations

The following examples show the relationship between VTAM pacing parameters and the initial pacing window sizes. Various network configurations are provided to illustrate a variety of pacing stages.

Same domain application program-to-application program session

For same domain application program-to-application program session pacing, there is one-stage adaptive pacing. [Figure 63 on page 225](#) shows the pacing stage for this example. See [Table 16 on page 225](#) and [Table 17 on page 226](#) for examples of the values used by VTAM to set the pacing-window values.




Figure 63. Same domain application program-to-application program session

Table 16 on page 225 shows sample definition statement coding and corresponding initial send windows when data flows from the PLU to the SLU.

Table 16. Same domain application program-to-application program session—PLU to SLU flow			
SRCVPAC MODEENT macro	VPACING APPL definition statement (SLU)	AUTH APPL definition statement (PLU)	Send window stage x
6	n/a	PACE	6
0	500	PACE	500
n/a	n/a	NVPACE	32767

Table 17 on page 226 shows sample definition statement coding and corresponding initial send windows when data flows from the SLU to the PLU.

<i>Table 17. Same domain application program-to-application program Session—SLU to PLU Flow</i>		
SSNDPAC MODEENT macro	VPACING APPL definition statement (PLU)	Send window stage x
0	n/a	32767
50	7 (default)	7
600	0	32767

Same domain application program-to-local device session

For same domain application program-to-local device session pacing there are two pacing stages. One stage, which is adaptive, exists between the application program and VTAM boundary function. Another pacing stage, which can be either fixed or adaptive, exists between the boundary function and the device LU. Figure 64 on page 226 shows the pacing stage for this example. See Table 18 on page 226 and Table 19 on page 227 for examples of the values used by VTAM to set the pacing-window values.

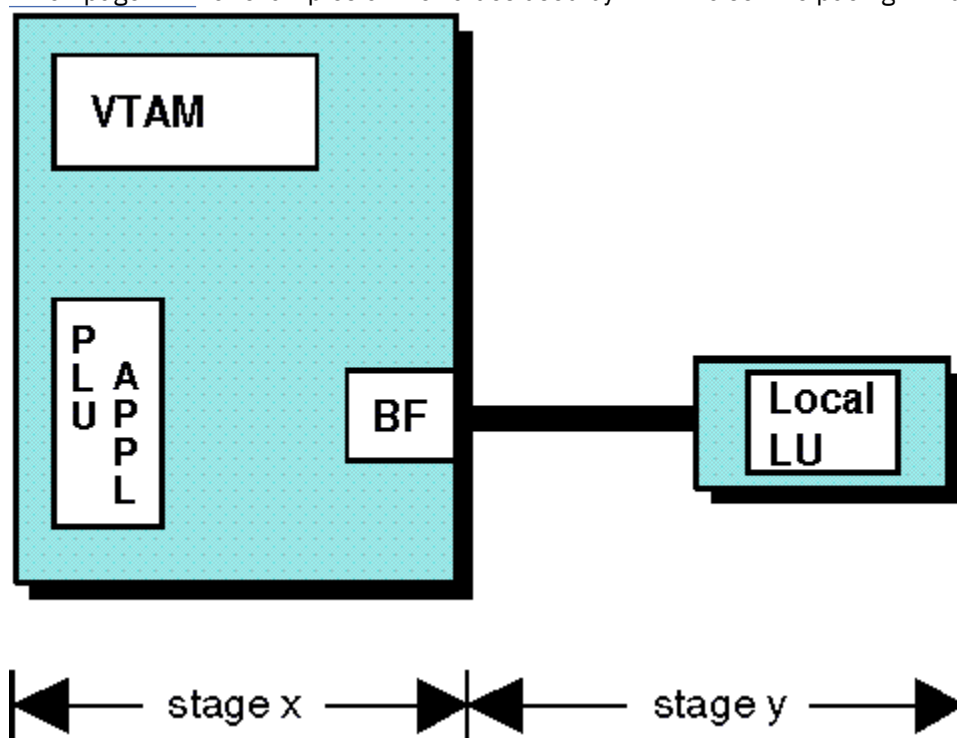


Figure 64. Same domain application program-to-local device session

Table 18 on page 226 shows sample definition statement coding and corresponding initial send windows when data flows from the PLU to the SLU.

<i>Table 18. Same domain application program-to-local device session—PLU to SLU flow</i>			
SRCVPAC MODEENT macro	VPACING definition statement (SLU)	Send window stage x	Send window stage y
6	n/a	6	6

Table 18. Same domain application program-to-local device session—PLU to SLU flow (continued)			
SRCVPAC MODEENT macro	VPACING definition statement (SLU)	Send window stage x	Send window stage y
0	500	500	500

Table 19 on page 227 shows sample definition statement coding and corresponding corresponding initial send windows when data flows from the SLU to the PLU.

Table 19. Same domain application program-to-local device session—SLU to PLU flow			
SSNDPAC MODEENT macro	VPACING APPL definition statement (PLU)	Send window stage x	Send Window Stage y
0	n/a	32767	32767
50	7 (default)	7	7
600	0	32767	32767

Application program-to-application program over APPN host-to-host connection

There are three pacing stages for application program-to-application program sessions over an APPN host-to-host connection. See Table 20 on page 228 and Table 21 on page 228 for examples of the values used by VTAM to set the pacing-window values. Figure 65 on page 227 show the pacing stages for this example.

Table 20 on page 228 shows sample definition statement coding and corresponding initial send windows when data flows from the PLU to the SLU.

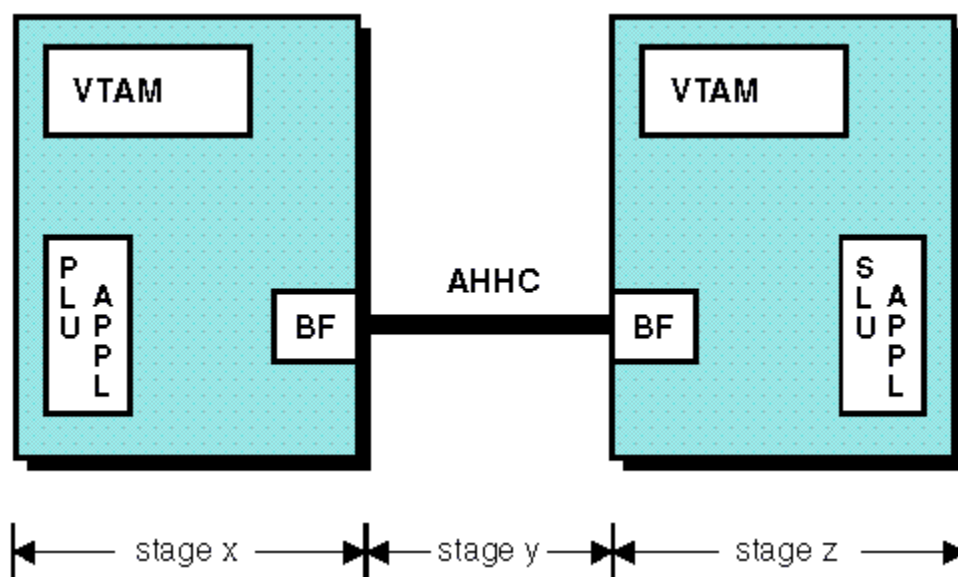


Figure 65. Application program-to-application program over APPN host-to-host connection

Table 20. Application program-to-application program over APPN host-to-host connection—PLU to SLU flow

SRCVPAC MODEENT macro	VPACING APPL definition statement (SLU)	Send window stage x	Send window stage y	Send window stage z
0	0	7	128	32767
0	7	7	7	7
50	n/a	50	50	50

Table 21 on page 228 shows sample definition statement coding and corresponding windows when data flows from the SLU to the PLU.

Table 21. Application program-to-application program over APPN host-to-host connection—SLU to PLU flow

SSNDPAC MODEENT macro	VPACING APPL definition statement (PLU)	Send window stage x	Send window stage y	Send window stage z
0	n/a	32767	128	7
50	7	7	7	7
50	0	32767	128	7

Application program-to-application program over CTCA connection

There is one pacing stage, which is adaptive, for application program-to-application program session over a CTCA connection. Figure 66 on page 229 shows the pacing stage for this example. See Table 22 on page 229 and Table 23 on page 229 for examples of the values used by VTAM to set the pacing-window values.

Table 22 on page 229 shows sample definition statement coding and corresponding windows when data flows from the PLU to the SLU.




Figure 66. Application program-to-application program over CTCA connection

Table 22. Application program-to-application program over CTCA connection—PLU to SLU flow			
SRCVPAC MODEENT macro	VPACING APPL definition statement (SLU)	AUTH operand (PLU)	Send window stage x
6	n/a	PACE	6
0	50	PACE	50
n/a	n/a	NVPACE	32767

Table 23 on page 229 shows sample definition statement coding and corresponding windows when data flows from the SLU to the PLU.

Table 23. Application program-to-application program over CTCA connection—SLU to PLU flow		
SSNDPAC MODEENT macro	VPACING APPL definition statement (PLU)	Send window stage x
0	n/a	32767
50	7 (default)	7
50	0	32767

Application program-to-local SNA over CTCA connection

There are two pacing stages for an application program-to-local SNA device session over a CTCA connection. Stage x uses adaptive pacing and stage y can use either adaptive or fixed pacing. Figure 67 on page 230 shows the pacing stages for this example. See Table 24 on page 230 and Table 25 on page 230 for examples of the values used by VTAM to set the pacing window values.




Figure 67. Application program-to-local SNA device over CTCA connection

Table 24 on page 230 shows sample definition statement coding and corresponding windows when data flows from the PLU to the SLU.

Table 24. Application program to local SNA device over CTCA connection—PLU to SLU flow						
SRCVPAC MODEENT macro	PSNDPAC MODEENT macro	PACING operand (SLU)	VPACING operand (SLU)	AUTH APPL definition statement (PLU)	Send window stage x	Send window stage y
0	n/a	n/a	4	NVPACE	95	4
0	15	n/a	4	PACE	15	4
5	0	6	n/a	PACE	6	5

Table 25 on page 230 shows sample definition statement coding and corresponding windows when data flows from the SLU to the PLU.

Table 25. Application program-to-local SNA device over CTCA connection—SLU to PLU flow			
SSNDPAC MODEENT macro	VPACING APPL definition statement (PLU)	Send window stage x	Send window stage y
9	2	2	9
0	7 (default)	7	32767
9	0	32767	9

Application program-to-local SNA device over AHHC connection

There are three pacing stages for application program-to-local SNA device session over an AHHC connection. Stage x and stage y use adaptive pacing and stage z uses either adaptive or fixed pacing. Figure 68 on page 231 shows the pacing stages for this example. See Table 26 on page 231 and Table 27 on page 231 for examples of the values used by VTAM to set the pacing window values.




Figure 68. Application program-to-local SNA device over AHHC connection

Table 26 on page 231 shows sample definition statement coding and corresponding windows when data flows from the PLU to the SLU.

Table 26. Application program-to-local SNA device over AHHC connection—PLU to the SLU						
SRCVPAC MODEENT macro	PSNDPAC MODEENT macro	VPACING operand (PLU)	PACING operand (PLU)	Send window stage x	Send window stage y	Send window stage z
5	0	6	n/a	6	5	5
5	0	6	n/a	6	5	5
0	15	n/a	4	15	4	4

Table 27 on page 231 shows sample definition statement coding and corresponding windows when data flows from the SLU to the PLU.

Table 27. Application program-to-local SNA device over AHHC connection—SLU to PLU flow					
AUTH APPL definition statement (PLU)	SSNDPAC MODEENT macro	VPACING APPL definition statement (PLU)	Send window stage x	Send window stage y	Send window stage z
PACE	0	7 (default)	7	128	32767
NVPACE	100	n/a	32767	128	100
PACE	9	8	8	9	9

Application program-to-application program with VR from intermediate host-to-SLU host

There are three pacing stages for application program-to-application program with VR from an intermediate host to the SLU host connection. All three pacing stages use adaptive pacing. Figure 69 on page 232 shows the pacing stages for this example. Stage x and stage y use adaptive pacing and stage z uses either adaptive or fixed pacing. Figure 69 on page 232 shows the pacing stages for this example.

See Table 28 on page 232 and Table 29 on page 232 for examples of the values used by VTAM to set the pacing window values. Figure 69 on page 232 shows the pacing stages for this example.




Figure 69. Application program-to-application program with VR from intermediate host-to-SLU host

Table 28 on page 232 shows sample definition statement coding and corresponding windows when data flows from the PLU to the SLU.

Table 28. Application program-to-application program with VR from intermediate host-to-SLU host—PLU to SLU flow

SRCVPAC MODEENT macro	VPACING APPL definition statement (SLU)	Send window stage x	Send window stage y	Send window stage z
0	0	7	128	32767
0	7	7	7	7
50	n/a	50	50	50

Table 29 on page 232 shows sample definition statement coding and corresponding windows when data flows from the SLU to the PLU.

Table 29. Application program-to-application program with VR from intermediate host-to-SLU host—SLU to PLU flow

SSNDPAC MODEENT macro	VPACING APPL definition statement (PLU)	VR window VRPWSxx	Send window stage x	Send window stage y	Send window stage z
0	n/a	(128,255)	32767	128	255
50	7	n/a	7	7	7
50	0	(128,255)	32767	128	255

Logon and logoff requests from dependent logical units

To request a session with an application program, an LU sends a logon request to VTAM specifying the application program name and, optionally, a logon mode name and some additional information. Logon and logoff requests can be field-formatted SNA Initiate and Terminate requests, which are built by the logical unit that sent them, or they can be (unformatted) character-coded commands, which are

translated by VTAM into a field-formatted SNA request. VTAM uses session-level USS tables and interpret tables to do the translation.

A session-level USS table can be used to convert the USS command into a field-formatted SNA request. If a character-coded command violates the USS command syntax rules (see the [z/OS Communications Server: SNA Resource Definition Reference](#)), an interpret table must be used for the conversion.

You can customize VTAM to accept other character-coded commands by coding an interpret table. The following forms of requests make use of the interpret table to determine the name of the application program to which the request is being made:

- A field-formatted SNA request received directly from a logical unit
- An unformatted character-coded command, converted to a formatted SNA request

For example, any of the above forms of a request can indicate a logon to APPL01, and the interpret table can change APPL01 into TESTAPL1.

You associate an interpret table to a logical unit using the LOGTAB operand on the LU definition statement. You associate a USS table to a logical unit using the USSTAB operand on the LU definition statement.

For a complete description of operation- and session-level USS tables, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Unformatted logon requests using mixed-case passwords

Mixed case passwords can be used by applications if an SAF-compliant security product (such as RACF®) has enabled this support. In some cases, the USS LOGON DATA parameter is used to send the password to the application. If a terminal user enters a mixed case password on the USS LOGON command and it is translated to uppercase by the translation table, the logon will fail if the target application expects to receive the password in mixed case.

The USS LOGON command is displayed on the terminal as it is typed. Therefore, the password is displayed. For additional security, you can inform the terminal user to discontinue entering the password as part of the USS LOGON process. Instead, the application should prompt the terminal user for the password in a non-displayed field. If mixed case passwords are used and the terminal user continues to enter the password as part of the USS LOGON command, the logon will fail when using TRANSLATE=YES (the default) on the USSPARM because the password has been translated to uppercase.

To continue to allow the terminal user to enter the password on the LOGON command, use one of the following methods to support mixed case passwords. With any of these methods, if the user ID is entered with the password, you must first verify whether the application supports translating the user ID to uppercase. A simple test is to enter the DATA portion of the USS LOGON within single quotation marks with the user ID specified in lowercase. USS will not translate data within single quotation marks and the quotation marks are removed before the data is passed to the application. If the logon fails, the application does not support translating the user ID to uppercase and a terminal user must enter the user ID in uppercase and the password in mixed case for the methods suggested.

- If the current USS translate table is used to set all characters to uppercase, add the TRANSLATE=NO operand to the USSPARM macro for the corresponding USSCMD to prevent translation of the specified DATA USSPARM containing the password to uppercase. See the [z/OS Communications Server: SNA Resource Definition Reference](#) for details.
- Because USS will not process data within single quotation marks against the translate table (the quotation marks are removed before the data is passed to the application), instruct terminal users to enter the DATA portion of the USS LOGON within single quotation marks.
- If the terminal user is specifying only APPLID and DATA on the USS LOGON command, an interpret table can be used. The interpret function passes all entered data to the application in untranslated format. Use REMOVE=Y on the LOGCHAR macro to remove the first non-blank string from the entered data. An interpret table should be set up to look for both an upper and a lowercase APPLID because the terminal user could enter the APPLID in lowercase.

Chapter 12. Network routing

Routing within a network is different for a subarea and an APPN network. In a subarea network, network paths are designed and coded by the user. In an APPN network, routing is a dynamic function between nodes.

Network routing and resource location for APPN nodes

VTAM uses topology and routing services (TRS) and directory services (DS) to monitor and maintain information about the topology and resources of the network.

TRS maintains information about nodes, transmission groups (TGs), and classes of service so that appropriate routes through an APPN network can be calculated.

TRS creates a network topology database, in which it maintains information about the current topology of the intermediate routing portion of an APPN network (network nodes and the transmission groups interconnecting them). This network topology database is replicated at every network node in an APPN network. Whenever a new network node or TG is added, or whenever the characteristics of an old resource change, a topology database update (TDU) that describes the current characteristics of the resource is generated and propagated to all network nodes.

When a network node first joins an APPN network (that is, establishes CP-CP sessions with other nodes), the other nodes send the joining node a copy of the network topology database. This is known as topology exchange. (A topology exchange that occurs when a network node first joins an APPN network is also called a full topology exchange or an initial topology exchange.) If the joining node later experiences planned or unplanned down time, it can remonitor the network topology by receiving all the topology updates that were exchanged since its topology database was last updated.

Directory services maintains a database that contains information about the location of LUs in the network. Entries are added to this database by resource definition, resource registration, and dynamically through searches. This means even if a session is requested to an LU whose location is not currently stored in the database, the location of this LU can be monitored, stored in the database, and session establishment will occur. Also, if a session is requested to an LU that cannot be located, that information is stored in the database if search reduction is enabled (see [“Improving VTAM performance using start options”](#) on page 24).

When a session is requested with an LU, VTAM uses its search capability and the directory services database to determine what node the resource is on. The search results give the location of the resource but not the route. VTAM dynamically determines the best route for the session using the APPN Class of Service requested for the session and the characteristics of the intermediate nodes and links stored in the topology database.

The defaults for the definition of the node and link characteristics and the VTAM ability to dynamically update the contents of the databases allow sessions to be established without requiring any definitions for resources outside the domain of that VTAM.

There are techniques you can use to improve the efficiency of resource location and session setup. These techniques include coding a central directory server and registering resources to a central directory server or a network node server. Resource registration occurs by default for some resources, but you can modify which resources are registered.

You can also checkpoint the topology and directory databases to ensure that information is retained during system outages. Checkpointing the directory database increases the chances that a resource can be located with a directed search instead of a broadcast search. Directed and broadcast searches are described in [“Types of searches”](#) on page 236. Checkpointing the topology database reduces the number of topology database updates that must flow at startup.

You may also need to adapt the Class of Service definitions, link and node characteristics to your network, such as coding some links as secure and requesting a secure Class of Service for certain sessions.

The following sections give information about searches, checkpointing, and APPN Class of Service.

Types of searches

To establish a session, the location of the partner resource must be known to VTAM. If the location of the partner resource is not statically defined, it can be monitored dynamically as a result of a directed, broadcast, or cross-subnetwork search. A directed search is used when VTAM has location information for the resource as the result of a previous search or if there is a central directory server in the network. A broadcast search is used when VTAM has no information about the location of the resource. A cross-subnetwork search is used to determine the location of a partner resource in another subnetwork.

A VTAM network operator can also initiate a search for a resource. Unlike searches that are performed for session initiation requests, an operator-initiated search does not update the directory services database when a resource is found.

Directed search

What is a directed search?

A directed search uses information stored in the directory services database of a network node server to direct the search to the location of the requested LU. A directed search is sent to the node recorded as the owner of the requested LU to verify the information. A directed search can be sent to a network node server from an end node and to an end node from a network node server. A directed search can also be sent from a network node to a central directory server when the network node does not have information about the location of the destination logical unit or following a failed search.

Who originates a directed search?

A directed search is originated by a network node server for the LUs it serves, by a central directory server as a result of a search request from a network node, or by an end node to its network node server.

When is a directed search needed?

A directed search is needed when a network node server or a central directory server receives a request for an LU for which it has a location stored in its database. The search is used to verify the information. A directed search is also used by a network node server to send a search request to a central directory server if the network node server has no information about the location of a requested resource. A directed search is used by an end node to request location information about a resource from the end node network node server.

Broadcast search

What is a broadcast search?

A broadcast search does not use database information about the location of a requested LU to propagate the search, because a directed search using database information has failed or there is no database information for the requested resource. Each network node receiving the search request sends the search to each of its adjacent network nodes; this allows the whole network to be searched. (VTAM end nodes do not allow searches for resources that are not registered. With other types of end nodes, broadcast searches include all domain end nodes that have requested that they be searched for the specific resource type.)

When the search reaches the network node serving the destination resource, that node sends back a positive reply to the first search request it receives. All subsequent search requests received by that network node as part of the same search are answered with a negative reply.

Who originates a broadcast search?

A broadcast search is initiated only by a network node server for the LUs it serves if there is no central directory server defined in the network. If there is a central directory server defined, only the central

directory server can issue a broadcast search. If a central directory server is active in your network, a network node server sends a directed search to the central directory server instead of issuing a broadcast search. Only if the directed search is unable to reach the central directory server does the network node server issue its own broadcast search.

When is a broadcast search needed?

A broadcast search is needed when there is no directory information giving the location of a requested LU or a directed search using existing directory information fails.

Broadcast searches in a network without a central directory server

In a network without a central directory server, broadcast searches are performed by any of the network nodes. These network nodes may be acting on behalf of LUs attached directly to them, LEN-attached LUs, or (when the network node is acting as a network node server) on behalf of LUs attached to end nodes that they serve. When the network node receives a request from an LU for a session with another LU, it checks for information in the database about the location of the partner LU. If it does not find the location information for the partner LU, the network node server of the requesting logical unit initiates a broadcast search.

The location information of the resource is not communicated to the other network nodes. If another network node server receives a request for the same resource and does not have the location information, it issues a broadcast search.

Broadcast searches in a network with a central directory server

In a network with a central directory server, broadcast searches are performed only by the central directory server. If a network node server does not know the location of a partner LU, it sends a directed search to the central directory server instead of issuing a broadcast search. The central directory server checks its database and issues a broadcast search only if it does not have the information. When the requested LU is located, the central directory server updates its directory and returns the information to the requesting network node. The network node then updates its database with the information and initiates the session with the partner LU. After a resource location is stored in the central directory database, no further broadcast searches are required, until the resource moves or is deactivated.

Cross-subnetwork search

What is a cross-subnetwork search?

A cross-subnetwork search is designed to find an LU outside the local APPN subnetwork. Because other APPN searches are sent between nodes that are in the same APPN subnetwork, this search type is employed at border nodes to forward searches for a requested LU to other APPN subnetworks. The target nodes of a cross-subnetwork search may be found in the directory services database, monitored dynamically, or explicitly defined by definition statements.

Who originates a cross-subnetwork search?

The first border node encountered as the search leaves the requester is responsible for originating a cross-subnetwork search. The border node forwards the request over locally managed subnetwork boundaries. Other nodes in nonnative subnetworks then perform their search logic to find the resource in their local subnetwork or across their subnetwork boundaries. These searches are performed sequentially until the target resource is located.

Is a cross-subnetwork search a broadcast search sent over subnetwork boundaries?

No. Cross-subnetwork search logic finds resources across subnetwork boundaries with minimum disturbance to other subnetworks. The sequential fashion in which the searches are sent allows the search to terminate as soon as the resource is located, without having to search every connected

subnetwork. In addition, routing information can be both defined or monitored dynamically, which directs the search toward the subnetwork that is most likely to own the resource.

Operator-initiated search

What is an operator-initiated search?

An operator-initiated search performs all valid searches for the node, including the following searches:

- Locally-attached subarea
- Subarea networks attached through other interchange nodes
- Broadcast
- Cross-subnetwork
- Local directory (directed search)

The operator-initiated search is started when a DISPLAY DIRECTRY, ID=*resource_name*, SCOPE=NSEARCH command is issued at a network node. The search returns all instances of the resource in the network. Unlike other searches, information returned by an operator-initiated search is not recorded in the directory services database.

When is an operator-initiated search needed?

The following are some examples of when you would use the operator-initiated search:

- During resource name administration, to verify that names chosen are unique within the network.
- To determine the owning control point of a resource.
- To determine whether duplicate names exist in the network that might result in session failures. If the information returned indicates that a naming conflict exists in the network, you can use the MODIFY RESOURCE command to change the registration of one of the resources.

Minimizing broadcast searches

You can control searches in the following ways:

- Define a central directory server to minimize the number of broadcast searches.
- Register resources to a directory database.
- Manage duplicate resource definition.

Search reduction also limits broadcast searches for unreachable viable resources in the network. See [“Improving VTAM performance using start options” on page 24](#) for additional information about search reduction.

Also, an installation-wide directory services management exit provides control over the extent to which VTAM conducts search requests in the network in which the directory services management exit is loaded. The directory services management exit is the interface between VTAM directory services and user-written code. VTAM conducts searches based on steps authorized by the user code. For more information about the directory services management exit, see [z/OS Communications Server: SNA Customization](#).

Central directory servers

A central directory server is a focal point for broadcast searches in the network. It uses broadcast and directed searches to locate resources. When a central directory server is active in the network, all network node servers issue directed searches to the central directory server instead of issuing a broadcast search.

What are the benefits of a central directory server?

A central directory server reduces the number of broadcast searches in your network. When the central directory server performs a broadcast search for a network node, it stores the location of the requested

LU in its database. If another network node server sends a directed search for the same resource to the central directory server, the central directory server does not have to issue another broadcast search because it already has the resource in its database.

Because a central directory server can receive search requests from many network nodes, its database grows faster and contains more resources than that of a network node server, which only receives the search requests of its served LUs. This centralization increases the chance that a resource will be found in the database and a broadcast search will not be required.

How does a central directory server receive a search request?

When a network node receives a search request, it checks its database for the resource. If it does not find the resource in its database, it sends the request to a central directory server if one exists in the network. If there is more than one central directory server in the network, the network node selects a central directory server based on one of the following possibilities:

- If the CDS selection function of the directory services management exit is active, the network node selects a central directory server from the list arranged by the exit. For information about how to use this function, see *z/OS Communications Server: SNA Customization*.
- If the CDSREFER start option value is 1 (the default), the network node sends the search request to the nearest (minimal weight route) central directory server in the network. For more information about the CDSREFER start option, see the [z/OS Communications Server: SNA Resource Definition Reference](#). If the CDS selection function of the directory services management exit is active, then the CDSREFER start option value is used only for ordering the list of central directory servers passed to the exit.

How does a central directory server handle a search request?

When the central directory server receives a search request, it checks its database for the resource. If it does not find the resource in its database and there are other central directory servers in the network, it sends the search to the other central directory servers only. The selection of alternate central directory servers can be customized by using the alternate CDS selection function of the directory services management exit. If the central directory server receives a positive reply from any of the other central directory servers, it verifies the information, updates its own database with the information, and notifies the originating network node of the location of the requested LU. If the central directory server receives negative replies from all the other central directory servers, it initiates a broadcast search. If the central directory server that received a search request locates the resource in its own database, it verifies the information with a directed search and sends a reply to the originating network node server with the location of the requested LU.

If the central directory server does not find the resource in its database and there are no other central directory servers in the network, it initiates a broadcast search.

What do you need to define to take advantage of a central directory server?

Use the CDSERVER start option to define VTAM as a central directory server. You do not need to define anything on the other network nodes. The other network nodes will find out about the existence of the central directory server through normal topology exchanges.

Registering resources

Resource registration places information about the location of resources in a directory services database. This registration reduces broadcast searches by ensuring that a resource will be found in the directory services database. Resources can be registered to a directory database on a network node server or to a central directory server, or both. By registering resources to a network node server, you ensure that the resources will be found during a search. By registering resources to a central directory server, you ensure that the resources will be found without requiring a broadcast search.

Resources are registered by using the REGISTER operand on the APPL, CDRSC, LOCAL, and LU definition statements. This operand can also specify that a resource is not registered. The defaults for registration change depending on the type of resource and whether the resource is located on an end node or a network node. [Table 30 on page 240](#) shows the default registration values for resources.

Table 30. Default registration values for resources

Major node	Definition statement	Default registration
Application program	APPL	Central directory server (REGISTER=CDSERVER)
CDRSC	CDRSC	No registration (REGISTER=NO)
Local Non-SNA	LOCAL	Network node server (REGISTER=NETSRVR)
LUGROUP	LU	Network node server (REGISTER=NETSRVR)
Local SNA Model NCP Switched	LU	Network node server for dependent LUs, LOCADDR > 0 (REGISTER=NETSRVR) No registration for independent LUs, LOCADDR = 0 (REGISTER=NO)

Dynamic registration for a resource can be initiated while VTAM is running by using the MODIFY RESOURCE command. For more information, see [z/OS Communications Server: SNA Operation](#).

Because the characteristics of an independent LU and of the node on which it resides are not recorded by VTAM, it is difficult to determine the appropriate level of registration for all cases. Therefore, independent LUs are not registered by default.

VTAM end nodes register their dependent LUs with their network node server by default. This includes all resources defined by LOCAL definition statements and LU definitions with a nonzero LOCADDR value. A resource on a VTAM end node that is not registered to the network node server will not be found during a search unless that resource has initiated a request through the network node server. Registering all end node resources ensures that they will be located by a search. You can also specify that these resources be registered to a central directory server.

Applications are registered to a central directory server by default, with the exception of TSO applications. For more information about default registration of TSO applications, see [“Resource registration in an APPN network”](#) on page 561.

Note: If a resource is never going to be a destination resource, there is no need to register it.

Resources on a network node can also be registered to a central directory server if one exists. By default, network nodes register only their applications to a central directory server.

Controlling which central directory server receives registration requests

If more than one central directory server exists in the network and the value for the CDSREFER start option is 1 (the default), the network node sends all central resource registration requests to the nearest (minimal-weight route) central directory server. You can change this value if you want to disperse directory information to other central directory servers in the network. For more information about this start option, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

You can alternatively specify which central directory servers receive registration requests by using the central resource registration CDS selection function of the directory services management exit. For information about how to use this function, see the [z/OS Communications Server: SNA Customization](#).

Controlling the timeliness and size of the directory database

By using the start options DIRTIME and DIRSIZE, you can control how long a resource remains in the directory services database and how many dynamic entries are maintained.

The DIRTIME start option specifies the amount of time a dynamic resource remains in the directory database without being used before it is deleted. The default value for this start option is eight days.

The DIRSIZE start option defines the maximum number of dynamically added resources that are stored in the database. After the number of dynamic resources reaches the maximum, new resources are added to the database in place of the resources that have gone the longest without being referenced.

Network search overhead caused by duplicate resource definitions

In an effort to facilitate recovery procedures, you might choose to define the same application programs or dependent LUs on more than one VTAM simultaneously. Because SNA does not support duplicate resource definition, your search logic must be able to locate the active instance of the resource first. If a nonactive instance is located first, the session usually fails to set up. However, when an APPN broadcast search is performed, you have little or no control over which instance is located first, and this can result in intermittent session failures.

To allow continued use of duplicate resource definitions, the APPN search logic in VTAM was modified to recognize nonactive resource search replies. The search will continue in the APPN network for the active instance of the specified resource. If an active instance is located, the reply from that instance is used and session setup continues. If an active instance is not located, the reply from a nonactive instance is used instead.

Although this additional search logic allows the continued use of duplicate resource definitions in VTAM, it does this by increasing the amount of network search traffic generated by VTAM. If your application program or dependent LUs are not in the active state, you might incur this overhead, even if you are not using duplicate resource definitions. To minimize or avoid this additional searching, you can use DUPDEFS start option. This start option lets you specify what types of resources (application programs, dependent LUs, or both) have duplicate definitions on this and other VTAMs.

Specifying DUPDEFS=ALL means that application programs and dependent LUs may be duplicated on this and other VTAMs. DUPDEFS=APPL means that only application program definitions may be duplicated. DUPDEFS=DEPLU means that only dependent LU definitions may be duplicated. In all cases, if the target resource is one that might have duplicate definitions and the resource is not in the ACTIVE state, VTAM returns a reply indicating that searching should continue to find an active instance of the resource. If the target resource is one that does not have duplicate definitions, the search reply returned by VTAM will indicate that searching should not continue (regardless of the current state of the resource). DUPDEFS=NONE means that no duplicate definitions exist on this and other VTAMs, which prevents the additional search logic from being performed.

Avoiding congestion

When search requests and search replies cannot be delivered immediately to a node, VTAM queues the data for delivery when possible. In situations where a node is not responding and searches are accumulating rapidly, VTAM storage can be excessively used. To prevent this overuse of VTAM storage, the MAXLOCAT start option can be specified.

The value of the MAXLOCAT operand is compared to the product of the number of queued searches and the amount of time between responses from the node. If the product is greater than the MAXLOCAT value, VTAM issues message IST1601I, which indicates that routing to this node is suspended and VTAM has stopped sending searches to the node until the congestion is reduced.

Even after VTAM suspends sending APPN search requests to an adjacent CP, VTAM continues to queue APPN search replies to the adjacent CP. If the number of queued requests plus the number of queued replies reaches 200% of MAXLOCAT and a minimum time interval has been exceeded, VTAM issues message IST2243I indicating that VTAM will initiate termination of the CP-CP session to the adjacent CP. VTAM ends the CP-CP session with sense code 80030004 and tries the CP-CP session again. This allows the storage related to the queued APPN searches to be cleaned up.

The minimum time interval is based on how long it has been since VTAM issued message IST1601I indicating routing to this node has been suspended. If IOPURGE is less than the default of 180 seconds,

then the minimum time interval is 180 seconds. Otherwise, the minimum time interval is the current IOPURGE value.

When coding a value for MAXLOCAT, consider the amount of VTAM storage available and the response time of nodes. For example, set the value high if there is plenty of available storage and a slow node exists.

Checkpointing of the TRS database and the directory database

When a network node is added to an APPN network, it must complete a full topology exchange with the other nodes in the network. In large APPN networks, full topology exchanges can involve the transfer of large amounts of data, affecting performance by slowing down the CP-CP session activation process.

When a node that has just joined the APPN network attempts to establish sessions, it has to monitor the location of all resources that it wishes to establish sessions with by issuing a broadcast search or a directed search to a central directory server.

By checkpointing the TRS and directory services databases, the topology and resource location information can be saved and used if VTAM is restarted.

Checkpointing saves the databases to the checkpoint data sets at the time either of the following conditions occurs:

- The VTAM operator issues a MODIFY CHKPT command.
- The VTAM operator issues a HALT or HALT QUICK command to terminate the VTAM network node.

This process of saving the databases is also called database hardening.

Note: The databases are not saved if the VTAM operator does the following actions:

- Issues a HALT CANCEL command to terminate the VTAM network node
- Cancels VTAM by canceling VTAM partition

When a database is checkpointed, it is written to a checkpoint data set. This data set is read by VTAM when it is restarted and the information is placed in the databases. If checkpointing is not used, both the topology and directory services databases have to be completely rebuilt. When checkpointing is used, only those changes to resources and topology since the last checkpoint need to be monitored.

To specify which of the checkpointed databases are to be loaded at startup time, you can use the INITDB start option. By default, both the topology and directory services databases are loaded. You can also specify topology only, directory services only, or neither.

APPN Class of Service

What Is APPN Class of Service?

APPN Class of Service defines the characteristics of a route for a session. Unlike subarea Class of Service, which defines the actual routes to be used for a session, APPN Class of Service defines the characteristics of the nodes and transmission groups (TGs) to be used for a session. VTAM chooses an APPN route by matching the required characteristics defined in the COS definition with the actual characteristics of the nodes and TGs in the topology database.

What information is specified in the APPN Class of Service?

The APPN Class of Service indicates the required characteristics of the nodes and lines in a route and the transmission priority of sessions. The possible line characteristics that can be specified include the security, speed (line capacity), propagation delay, and cost (both per unit of time and per byte). There are also three user-defined parameters that can be used to specify additional characteristics. The node characteristics that can be specified are the level of congestion allowed when selecting a node for routing and the desirability of the node for intermediate routing. The desirability of using a node for intermediate routing is called route resistance.

The characteristics are specified in ranges, with a low weight assigned to the most desirable set of ranges and a high weight assigned to the least desirable set of ranges. The first row in the default Class of Service has a very low weight and has the most restrictive set of ranges. The next set of ranges is less restrictive but has a higher weight. (These sets of ranges are called LINEROWs and NODEROWs.) The last set of ranges has the widest range of acceptable characteristics but has a very high weight. By coding the rows in this manner, the best resources (lowest weight and best characteristics) are chosen when determining the route.

Do you have to code anything to make APPN Class of Service work?

IBM supplies seven default APPN Class of Service definitions. The TG and node characteristics also default. For APPN-only sessions, VTAM locates an available route without any additional coding.

What are the IBM-supplied default classes of service?

Descriptions of the seven IBM-supplied APPN Classes of Service follow. These classes define seven generic types of session traffic. See the [z/OS Communications Server: SNA Resource Definition Reference](#) for samples.

#BATCH

An APPN COS for LU-LU sessions that specifies a general batch-oriented COS that uses low transmission priority, and for which high bandwidth and low cost are considered more important than short delay.

#BATCHSC

An APPN COS for LU-LU sessions that specifies a general batch-oriented COS that uses low transmission priority, and for which high bandwidth and low cost are considered more important than short delay. A minimum security level is required.

#CONNECT

An APPN COS for LU-LU sessions that provides connectivity at medium transmission priority.

CPSVCMG

An APPN COS for CP-CP sessions that is used for network flows. It provides connectivity at network transmission priority.

#INTER

An APPN COS for LU-LU sessions that specifies a general, interactive COS that uses high-transmission priority, and for which short delay is considered more important than high bandwidth and lost cost.

#INTERSC

An APPN COS for an LU-LU session that specifies a general, interactive COS that uses high-transmission priority, and for which short delay is considered more important than high bandwidth and lost cost. A minimal security level is required.

SNASVCMG

An APPN COS for LU-LU CNOS sessions that provides connectivity at network transmission priority.

Three sets of IBM-supplied classes of service are available. Each set contains the seven classes described previously. The three sets are:

- COSAPPN

The definitions in the COSAPPN definition set are made up of 8-row LINEROW entries and 8-row NODEROW entries for all classes of service and are appropriate for most sessions. The LINEROW values for the #CONNECT class of service in COSAPPN are shown in [Table 31 on page 244](#).

<i>Table 31. COSAPPN #CONNECT class of service LINEROW values</i>							
LINEROW	CAPACITY	COSTBYTE	COSTTIME	PDELAY	SECURITY	UPARM1 UPARM2 UPARM3	WEIGHT
1	4M MAXIMUM	0 0	0 0	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	30
2	56000 MAXIMUM	0 0	0 0	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	60
3	19200 MAXIMUM	0 0	0 0	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	90
4	9600 MAXIMUM	0 0	0 0	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	120
5	19200 MAXIMUM	0 0	0 0	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	150
6	9600 MAXIMUM	0 128	0 128	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	180
7	4800 MAXIMUM	0 196	0 196	MINIMUM MAXIMUM	UNSECURE MAXIMUM	0 255	210
8	MINIMUM MAXIMUM	0 255	0 255	MINIMUM MAXIMUM	UNSECURE MAXIMUM	0 255	240

- ISTACST2

The definitions in the ISTACST2 definition set are made up of 12-row LINEROW entries and 8-row NODEROW entries for all classes of service except CPSVCMG and SNASVCMG, which are made up of 8-row LINEROW entries and 8-row NODEROW entries. These 12-row LINEROW entries enable z/OS Communication Server to select an optimal route for a session. The greater number of entries provides optimal results when multiple connections with different TG characteristics (for example, channel-to-channel) are used in the network.

The LINEROW values for the #CONNECT class of service in ISTACST2 are shown in [Table 32 on page 244](#).

<i>Table 32. ISTACST2 #CONNECT class of service LINEROW values</i>							
LINEROW	CAPACITY	COSTBYTE	COSTTIME	PDELAY	SECURITY	UPARM1 UPARM2 UPARM3	WEIGHT
1	100M MAXIMUM	0 0	0 64	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	20
2	10M MAXIMUM	0 0	0 64	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	40

Table 32. ISTACST2 #CONNECT class of service LINEROW values (continued)							
LINEROW	CAPACITY	COSTBYTE	COSTTIME	PDELAY	SECURITY	UPARM1 UPARM2 UPARM3	WEIGHT
3	4M MAXIMUM	0 0	0 64	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	60
4	10M MAXIMUM	0 0	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	80
5	56000 MAXIMUM	0 0	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	100
6	29K MAXIMUM	0 0	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	120
7	100M MAXIMUM	0 128	0 128	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	140
8	10M MAXIMUM	0 128	0 128	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	160
9	56000 MAXIMUM	0 128	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	180
10	29K MAXIMUM	0 128	0 128	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	200
11	9600 MAXIMUM	0 0	0 64	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	220
12	MINIMUM MAXIMUM	0 255	0 255	MINIMUM MAXIMUM	UNSECURE MAXIMUM	0 255	240

- ISTACST3

The definitions in the ISTACST3 definition set are made up of 12-row LINEROW entries and 8-row NODEROW entries for all seven classes of service. The definitions enable z/OS Communications Server to select an optimal route for a session when connections in the network include connections with high speed link characteristics such as FICON®, Gigabit Ethernet, and HiperSockets.

The LINEROW values for the #CONNECT class of service in ISTACST3 are shown in [Table 33 on page 245](#).

Table 33. ISTACST3 #CONNECT class of service LINEROW values							
LINEROW	CAPACITY	COSTBYTE	COSTTIME	PDELAY	SECURITY	UPARM1 UPARM2 UPARM3	WEIGHT
1	10G MAXIMUM	0 0	0 64	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	20

<i>Table 33. ISTACST3 #CONNECT class of service LINEROW values (continued)</i>							
LINEROW	CAPACITY	COSTBYTE	COSTTIME	PDELAY	SECURITY	UPARM1 UPARM2 UPARM3	WEIGHT
2	1G MAXIMUM	0 0	0 64	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	40
3	100M MAXIMUM	0 0	0 64	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	60
4	1G MAXIMUM	0 0	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	80
5	25M MAXIMUM	0 0	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	100
6	1M MAXIMUM	0 0	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	120
7	10G MAXIMUM	0 128	0 128	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	140
8	1G MAXIMUM	0 128	0 128	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	160
9	16M MAXIMUM	0 128	0 64	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	180
10	1M MAXIMUM	0 128	0 128	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	200
11	64K MAXIMUM	0 0	0 64	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	220
12	MINIMUM MAXIMUM	0 255	0 255	MINIMUM MAXIMUM	UNSECURE MAXIMUM	0 255	240

The definition sets COSAPPN, ISTACST2, and ISTACST3 are shipped in SYS1.ASAMPLIB. To use one of these sets, you must copy it into SYS1.VTAMLST when you install z/OS Communication Server, and then activate the member in which the definitions reside. You can copy multiple sets of definitions into SYS1.VTAMLST, but you can have only one set active at any time.

The COSAPPN definition set is automatically activated when z/OS Communication Server is initialized. If you choose to use ISTACST2 or ISTACST3, use the VARY ACT command to activate it, or place the ISTACST2 or ISTACST3 member in the configuration list to automatically activate it when z/OS Communication Server is initialized. You can rename the IBM-supplied sets of definitions so that ISTACST2 or ISTACST3 is named COSAPPN, and COSAPPN is either not used or is renamed to something else. Renaming the definition sets in this way causes the set of definitions with 12-row LINEROW entries to be automatically activated when z/OS Communication Server is initialized.

Guideline: If you activate a set of definitions with 12-row LINEROW entries, you should have the same set of definitions with 12-row LINEROW entries activated on each network node in the network for optimal routing in networks that include high-speed connections.

Not all HPR APPN products support COS definitions with 12-row LINEROW entries. This could affect your ability to optimally use high speed connections among the nodes in your network. Consult technical representatives for the HPR APPN products in your network to determine if those products support COS definitions with 12-row LINEROW entries.

If you use COS definitions with 12-row LINEROW entries, routes selected for sessions could be different than those selected when you use COS definitions with 8-row LINEROW entries.

What do you get if you accept all the defaults?

The name of the APPN Class of Service definition associated with the mode is specified by using the APPNCOS operand on the MODEENT macroinstruction. Values are supplied for the APPNCOS operands in the IBM-supplied logon mode table (ISTINCLM). The value for APPNCOS in the IBM-supplied logon mode table is coded with the APPN Class of Service that most closely matches the type of session the logon mode entry is intended for. For example, if the logon mode name is coded as BATCH, BAT13790, BAT23790, or #BATCH, the APPNCOS operand value is #BATCH³. For the complete list of logon mode names and APPNCOS values, see [Table 34 on page 247](#).

Table 34. ISTINCLM APPNCOS values based on LOGMODE operand values	
IBM-supplied mode name	APPNCOS value
CPSVCMG	CPSVCMG
CPSVRMGR	SNASVCMG
SNASVCMG	SNASVCMG
#BATCH BATCH BAT13790 BAT23790	#BATCH
#BATCHSC	#BATCHSC
#INTER INTERACT INTRACT INTRUSER	#INTER
#INTERSC	#INTERSC
DSILGMOD DSC4K D32901 EMUDPCX SCS xxx3270x xxx3277x xxx3278x xxx3790x etc.	#CONNECT
ISTCOSDF	#CONNECT

³ # represents X'7B'. You should use whichever character on your system equates to X'7B'.

Table 34. ISTINCLM APPNCOS values based on LOGMODE operand values (continued)	
IBM-supplied mode name	APPNCOS value
IBMRDB	#CONNECT

What is the APPN Class of Service default in a user-coded logon mode table?

If the APPNCOS operand is not coded on the MODEENT macroinstruction, VTAM checks the COS operand value.

- If the COS operand value is coded, VTAM checks for an APPNCOS definition by the same name as the COS name. If VTAM does not find an APPNCOS definition by the same name as the COS, it fails the session request, unless the APPNCOS start option is defined with a substitute APPN Class of Service name.
- If the COS operand is not coded, VTAM defaults APPNCOS to #CONNECT, regardless of the setting for the APPNCOS start option.

How do I add APPN Class of Service to my logon mode table?

You can add APPN Class of Service to your user-coded logon mode table in one of three ways:

- You can create APPNTOSA and SATOAPPN tables to allow VTAM to map Class of Service between APPN and subarea networks, without having to change existing logon mode tables.
- You can code the APPNCOS operand for each logon mode entry in the table.
- You can create new APPNCOS definitions with the same name as the subarea COS.

To create APPNCOS definitions with the same name as the subarea COS, you can copy the default definition for APPN Class of Service that best fits that logon mode and change the name to match the COS operand coded on that logon mode.

How does z/OS Communication Server use the Class of Service to choose a route?

z/OS Communication Server chooses a route by comparing the actual characteristics of the available nodes and TGs to the allowed characteristic ranges specified in the requested COS. For each APPN COS entry, there are 1–12 LINEROW operands and 1–8 NODEROW operands. These operands give up to 12 acceptable sets of characteristics for the lines and up to 8 for the nodes in each COS. The WEIGHT parameter on the NODEROW and LINEROW operands is coded to indicate the desirability of that set of characteristics. The lower the value of the WEIGHT parameter, the higher the desirability of a node or TG that fits that set of characteristics.

Tip: The IBM-supplied COS definitions in COSAPPN are made up of 8-row LINEROW and NODEROW entries for all classes of service. The IBM-supplied COS definitions in ISTACST2 are made up of 12-row LINEROW entries for all classes of service except CPSVCMG and SNASVCMG, (which are made up of 8-row LINEROW entries), and 8-row NODEROW entries for all classes of service. The IBM-supplied COS definitions in ISTACST3 are made up of 12-row LINEROW entries and 8-row NODEROW entries for all classes of service. See [Class of Service](#) for more information.

For each of the LINEROW and NODEROW operands, an acceptable range is coded. For each potential line or node in a route, z/OS Communication Server compares the node or TG characteristics to the LINEROW or NODEROW operand values in the requested Class of Service. For an entire route to be acceptable for a given Class of Service, all nodes and TGs for that route must be acceptable for the requested Class of Service.

If more than one route fits the required characteristics of the Class of Service, z/OS Communication Server chooses the one with the least total weight, and therefore the most desirable characteristics. If no

route matches the required characteristics of the Class of Service, the session fails (for example, if the Class of Service requires a secure route and VTAM cannot find a complete route using only secure TGs).

Influencing session routing

The RESUSAGE start option is used to affect the number of times a given resource is used in creating a route. RESUSAGE specifies the number of times a resource can be used before it is considered overused. When a resource is overused, any routing trees using it are reconstructed in an attempt to use alternate resources.

Example of weight computation

The following example illustrates how a class of service is used to assign a weight to an element of a route. This example illustrates what weight is assigned to a TG with default characteristics using the #CONNECT class of service. Default characteristics of a TG are shown in Table 35 on page 249. Default LINEROW values for the #CONNECT Class of Service are shown in Table 36 on page 249.

Table 35. Default TG characteristics	
TG number	TG1
CAPACITY	8K
COSTBYTE	0
COSTTIME	0
PDELAY	TERRESTR
SECURITY	UNSECURE
UPARM1	128
UPARM2	128
UPARM3	128

Table 36. #CONNECT Class of Service LINEROW values							
LINEROW	CAPACITY	COSTBYTE	COSTTIME	PDELAY	SECURITY	UPARM1 UPARM2 UPARM3	Weight
1	4M MAXIMUM	0 0	0 0	MINIMUM NEGLIGIB	UNSECURE MAXIMUM	0 255	30
2	56000 MAXIMUM	0 0	0 0	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	60
3	19200 MAXIMUM	0 0	0 0	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	90
4	9600 MAXIMUM	0 0	0 0	MINIMUM TERRESTR	UNSECURE MAXIMUM	0 255	120
5	19200 MAXIMUM	0 0	0 0	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	150

Table 36. #CONNECT Class of Service LINEROW values (continued)							
LINEROW	CAPACITY	COSTBYTE	COSTTIME	PDELAY	SECURITY	UPARM1 UPARM2 UPARM3	Weight
6	9600 MAXIMUM	0 128	0 128	MINIMUM PACKET	UNSECURE MAXIMUM	0 255	180
7	4800 MAXIMUM	0 196	0 196	MINIMUM MAXIMUM	UNSECURE MAXIMUM	0 255	210
8	MINIMUM MAXIMUM	0 255	0 255	MINIMUM MAXIMUM	UNSECURE MAXIMUM	0 255	240

The values shown in [Table 36 on page 249](#) for capacity, cost per byte, cost per unit of time, propagation delay, security level, and the user-defined characteristics (UPARM1, UPARM2, and UPARM3) represent ranges, with the top value in a LINEROW representing the minimum value and the bottom value in a LINEROW representing the maximum value. In comparing these values with the TG characteristics shown in [Table 35 on page 249](#), the first match occurs at row 7.

The cost per byte, cost per unit of time, security, and user-defined characteristics are acceptable for all rows. However, the possible LINEROW values for capacity (number of bits per second that the link can transmit data) are MINIMUM, 600, 1200, 2400, 4800, 9600, 14400, 19200, 48000, 56000, 64000, 4M, 16M, and MAXIMUM, and effective capacity is not within the range for rows 1 through 6. The possible values of propagation delay (from lowest to highest) are:

- MINIMUM
- NEGLIGIB
- TERRESTR
- PACKET
- LONG
- MAXIMUM

Propagation delay is not within the range for the first row.

So, the default values for a TG match both rows 7 and 8 of the #CONNECT Class of Service. Row 7 has a weight of 210 and this weight is assigned to this TG for the route being calculated. Row 8 has a weight of 240, but even if it had a lesser weight it would not be used. (It is recommended that you code your rows from least weight to highest weight to ensure that the least weight row is used in the event of multiple matches.)

This TG is used if no other TG is available. If another TG is available, it is used if its weight is less than 210. In the event two or more TGs are equally weighted, one is chosen at random.

If this TG is being considered for part of a route that involves nodes and other TGs, the weight of this TG would be added to the weights of the other TGs and nodes, to produce a total weight for the route.

If the origin or destination of this TG is a virtual node, the weight of the TG is divided by two and the calculated value is used in route calculation.

For an example of how to use the node and TG characteristics to force a route, see [Appendix H, "Forcing an APPN route in a VTAM network," on page 613](#).

APPN network routing through a composite network node (CNN)

In a configuration containing a composite network node (CNN), it is possible that the optimal route computation does not occur. Route computation is performed by the network node serving the origin

logical unit (OLU). For example, in [Figure 70 on page 251](#), a route from LU to APPL is computed that uses TG1-TG4. TG1-TG3 would be a better route.

Composite Network Node




Figure 70. Routing example through a CNN node

VTAM in the composite network node (CNN) node affects APPN routing by:

- Selecting an optimal route through the CNN

When computing a route and choosing among TGs of equal weight, VTAM searches for matching subarea numbers for the entry and exit TGs to a CNN (TG1 and TG3 in this example, because they share NCP1A). If TGs with matching subarea numbers are found, VTAM selects that route. In situations where no match is found (for example, if TG3 did not exist), the original route is used (either TG1–TG4 or TG1–TG2).

- Changing the route during BIND processing

The CNN node affects routing during BIND processing. When a BIND reaches the CNN node, the next hop in the RSCV contained in the BIND may be changed by VTAM in the CNN node when the entry and exit TGs use two different subareas. VTAM changes the route during BIND processing when the following conditions are met:

- The subarea number of the new exit TG is the same as the subarea number of the entry TG.
- TG characteristics are the same as the original TG.

[Figure 71 on page 252](#) shows an example of VTAM changing the selection of a route. In [Figure 71 on page 252](#) mesh connectivity exists between VTAMA, VTAMB, NCPA, NCPB, and NCPC. In selecting a path between EN1 and NN2, VTAMA determines that using TG2-TG3 is the optimal route. Because NN2 does not use subarea numbers in route calculation, it might choose TG2-TG4 when selecting a route from NN2 to EN1.




Figure 71. CNN route calculation example

Figure 72 on page 252 shows another example of route calculation through a CNN node. VTAMA, the dependent LU server for ENA, calculates a route of TG5-TG6.




Figure 72. Composite network node route calculation example

Figure 73 on page 253 is an example of the CNN node changing the route when receiving a BIND request. A BIND request sent from an LU located on NN1 to an LU on NN2 contains TG1-TG2 in the RSCV. VTAMA

changes the RSVC to use TG1-TG3. By rerouting the BIND to a TG with the same subarea as the entry TG, VTAM prevents unnecessary hops within the CNN.

Composite Network Node




Figure 73. Composite network node route BIND reroute example

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. For example, referring to Figure 73 on page 253, if NN2 calculates TG2-TG1 as the route to go from NN2 to NN1, when the BIND arrives at VTAMA it cannot be changed to use TG3. It is too late in the process for the CNN node to change the route. However, if CNNRTMSG=NOSUPP is in effect, VTAMA issues the IST1774I message group to indicate that an optimal route through the CNN does exist but was not chosen during session activation.

Note that all the optimizations concern *equally* weighted TGs. For this reason, making the characteristics of TGs between two identical nodes allows VTAM more flexibility in optimizing paths through CNNs.

Using the SAMAP table

The subarea mapping (SAMAP) table is used to supply additional NCP connectivity information to APPN topology and routing services (TRS), which can then use that information in an effort to optimize routes through the CNN.

In configurations with remote NCP, the relationship between the local and remote NCP is not known to APPN TRS. Therefore, the subarea number mapping process (described in the previous section) is not adequate. The SAMAP table can be used to specify local-to-remote NCP connectivity for use by APPN TRS during route computation.

The SAMAP table can also be used to map the subarea number for a gateway NCP to the Net ID for the SNI-connected network that will be accessed through that gateway NCP.

Subarea number mapping

While APPN topology and routing services is not normally aware of the subarea connectivity within the CNN, the SAMAP table can be used to provide NCP-to-NCP mapping information that can be used to improve the intra-CNN routing in many cases.

Figure 74 on page 254 illustrates a typical session path scenario.




Figure 74. Typical CNN session path

In the above scenario, each of the three NCPs (subarea 10, 14, and 16) are owned by the NN1 composite network node and connect to each of the two hosts (NN1 and EN2). Remote NCP SA20 is connected only to NCP/SA10. When LUs connected to SA20 attempt to log on to applications on EN2, the best route (through the link to NCP/SA10) is not always selected. [Figure 75 on page 254](#) shows an APPN view of the configuration shown in [Figure 74 on page 254](#).




Figure 75. APPN view of CNN session path

In [Figure 75 on page 254](#), three TGs are shown between NN1 and EN2, and the DLU is owned by SA20. If the SA10 TG is selected when VTAM calculates the route, the session is set up. If either 14 or 16 is selected, the resulting path passing through two NCPs is less efficient. This route might fail if there is no VR defined from NCP/SA14 or NCP/SA16 to NCP/SA20.

The SAMAP table provides a way to define to APPN TRS the relationship between the subarea components of the CNN. In this way, the TRS can use that information to determine the best links to use for sessions entering or leaving the CNN. The two subarea numbers coded in the SAMAP table represent two subarea components in the CNN that are connected to each other. The definition of this relationship is most important for CNNs that have remote NCPs.

[Figure 76 on page 254](#) illustrates a configuration where the SAMAP table would assist in determining optimal CNN routes.

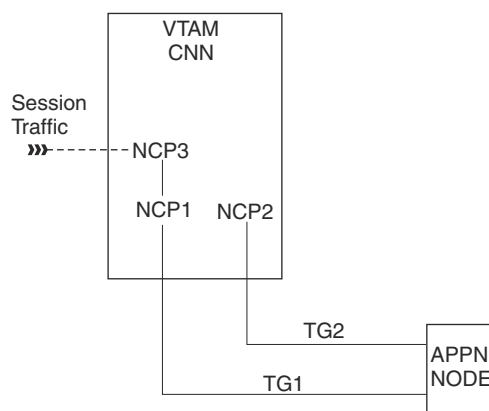


Figure 76. SAMAP session routing concept

In [Figure 76 on page 254](#), the VTAM CNN owns NCP1, NCP2, and NCP3. TG1 and TG2 have identical route weights based on their TG characteristics and the APPN COS being used. However, NCP3 is connected directly to NCP1 (not NCP2). At session setup, it is better to select TG1 instead of TG2. In this way, session traffic will traverse NCP3–NCP1 within the CNN. If the session was set up with TG2, session traffic would have to traverse the path NCP3–VTAM–NCP2 within the CNN to reach the next APPN node.

You will need to code SA03 MAPSTO SUBAREA=SA01 on the SAMAP table to indicate a direct connection exists between subarea 3 and subarea 1. For sessions to LUs attached to NCP3, TRS will select TG1 based on the SAMAP entry (assuming both TG1 and TG2 are operational and have the exact same route weight). If you do not use the SAMAP entry, TRS randomly selects either TG during session setup.

Note: The MAPSTO statement specifies that connectivity exists between the two subareas and is not directional in nature. In other words, SA03 MAPSTO SUBAREA=SA01 has the exact same meaning as SA01 MAPSTO SUBAREA=SA03. Only one of those statements needs to be included in the table.

The SAMAP table reflects the NCP-to-NCP connectivity within a CNN environment. Within a single SAMAP table, it is permissible to specify the NCP-to-NCP connectivity for all of the CNNs within the network, as long as the subarea numbers are all unique across all of the CNNs within the APPN subnetwork.

Net ID mapping

When computing a route to a partner in an SNI-connected network, APPN TRS might not always be able to choose the optimal NCPs. This is because the APPN portion of the route might be calculated before the address assignment for the gateway NCP. You can use the SAMAP table to inform TRS that a particular gateway NCP should be assumed when computing routes to a given Net ID.

To provide this information, use the Net ID form of the SAMAP entry:

```
SA03 MAPSTO NETID=NETX
```

This informs TRS that when computing routes to a partner in NETX, the gateway NCP with subarea number 03 will be used. Then, TRS can use that information during TG selection for the CNN.

Note: The SAMAP table on a given host can contain both types of SAMAP entries, and this is often appropriate for VTAMs that own SNI connections.

So, a SAMAP table can contain:

- SAxx MAPSTO SUBAREA=SAyy definitions to map all of the remote NCP to the local NCPs, and the gateway NCPs to local NCPs that have APPN connections.
- SAxx MAPSTO NETID=netid definitions to map the gateway NCP subarea numbers to the Net IDs of their SNI-partners.

Consider the example shown in [Figure 77 on page 255](#).

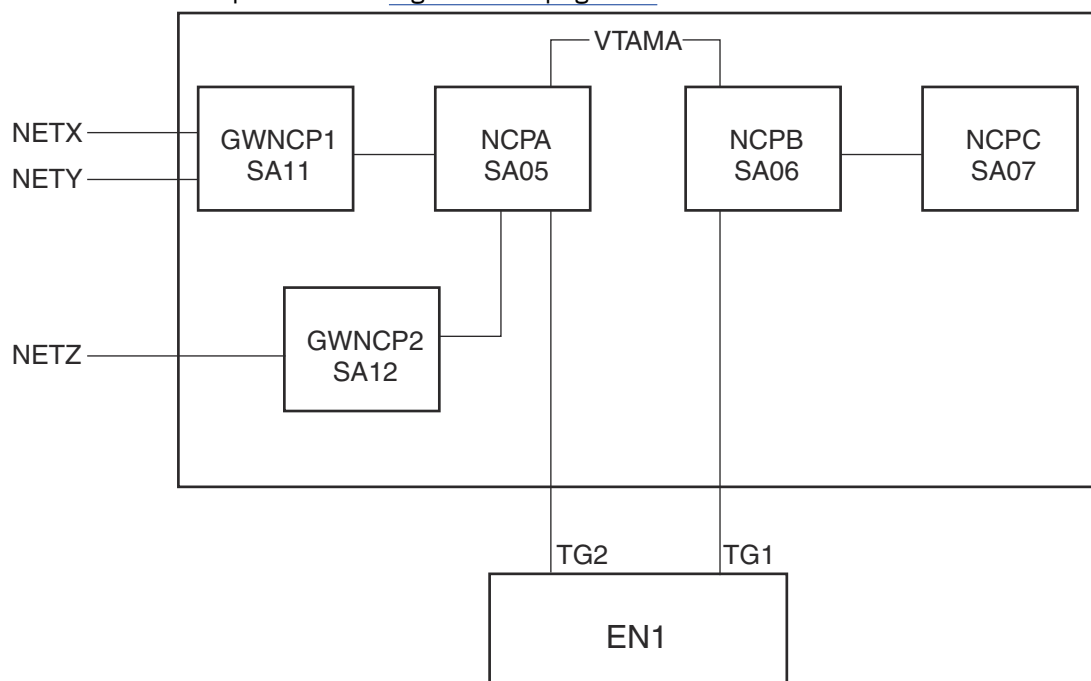


Figure 77. SAMAP example

In this case, the SAMAP table should reflect the following conditions:

- If the partner is in NETX or NETY, GWNCP1 (subarea 11) will be used.
- If the partner is in NETZ, GWNCP2 (subarea 12) will be used.
- NCPA (subarea 05) is connected to GWNCP1 and GWNCP2. This information allows TRS to prefer TG2 (assuming TG1 and TG2 are equally weighted) for sessions going through GWNCP1 and GWNCP2.
- NCPB (subarea 07) is connected to NCPA (subarea 06). This information allows TRS to prefer TG1 for sessions to dependent LUs attached to NCPB.

The corresponding SAMAP table would be as follows:

```
SAMAP1 VBUILD TYPE=SAMAP
SA11    MAPSTO SUBAREA=SA05
SA12    MAPSTO SUBAREA=SA05
SA07    MAPSTO SUBAREA=SA06
SA11    MAPSTO NETID=NETX
SA12    MAPSTO NETID=NETZ
SA11    MAPSTO NETID=NETY
```

The preceding example shows the subarea-to-subarea mappings together in a group, preceding the Net ID mappings. However, this sequence is only illustrative, and it is not required.

Network routing for subarea nodes

You need to design routes between:

- VTAM and any NCPs it activates
- NCPs communicating over a subarea connection (which includes NCPs within a composite network node)
- VTAMs communicating over a subarea connection

Communication between two network accessible units (NAUs) in different subareas requires a definition of at least one route connecting the subareas. This definition includes a physical and logical path between the two.

Physical paths

The physical path between two subarea nodes is an explicit route.

Explicit route (ER)

An explicit route is an ordered set of subarea nodes and transmission groups along a path between communicating subarea nodes, including:

- The endpoint subareas
- Any subareas between the endpoint subareas
- The transmission group used to connect each subarea pair along the route

In connecting one endpoint subarea to another, an explicit route cannot include the same subarea more than once.

Forward explicit route

Explicit routes originating in the host are forward explicit routes and are numbered 0 through 15 (0 through 7 with releases before NCP Version 4 Release 3.1 or Version 5 Release 2.1 or if ERLIMIT is set to 8).

Reverse explicit route

Reverse explicit routes terminate in the host and must use the same set of subarea nodes and transmission groups as their corresponding forward explicit route. They are also numbered 0 through 15 (0 through 7 with releases before NCP Version 4 Release 3.1 or Version 5 Release 2.1 or if ERLIMIT

is set to 8), but they do not have to have the same explicit route number as the corresponding forward explicit route.

Transmission group (TG)

A transmission group consists of a single SDLC link or data channel between two subarea nodes, or of a logical combination of up to eight parallel SDLC links between communication controller subarea nodes. Transmission groups are numbered 1 through 255.

Transmission groups appear as one link to the path control network, even if the transmission group contains parallel SDLC links. Data flows across a specific transmission group during a session, but not across a specific link in a transmission group. The specific link used within the transmission group is determined dynamically depending on which link is available at the time.

Transmission groups can also exist in parallel, creating multiple explicit routes between two subareas. For instance, parallel SDLC links between adjacent communication controller subarea nodes can be divided into multiple transmission groups. A maximum of 16 parallel transmission groups are allowed between adjacent subarea nodes.

Parallel links in one or more transmission groups can be active and carry data simultaneously. Each parallel link is controlled (activated and deactivated) independently of the others.

Routing tables

An entire explicit route is not defined to each subarea node along a path; instead, each subarea node has its own routing table that contains only the information that path control needs to forward data to the next subarea along the route (the adjacent subarea). Given the explicit route number being used and the destination subarea, the table provides the adjacent subarea to which the data is forwarded and the transmission group over which the data is forwarded.

Peripheral link

A peripheral link is the portion of the route between a subarea node and a peripheral node. A peripheral node uses local addresses for routing and requires boundary function assistance from an adjacent subarea node to communicate with a nonadjacent subarea node.

In the sample network configuration shown in Figure 78 on page 257, HOSTA and peripheral node T1 can communicate over two explicit routes, one in each direction. Both explicit routes have been numbered 0 and include HOSTA, NCPA1, NCPA2, and transmission groups 1 and 14. A data channel assigned to TG1 connects HOSTA with NCPA1. An SDLC link assigned to TG14 connects NCPA1 and NCPA2. Each subarea node contains a routing table to route data between subareas along ER0.

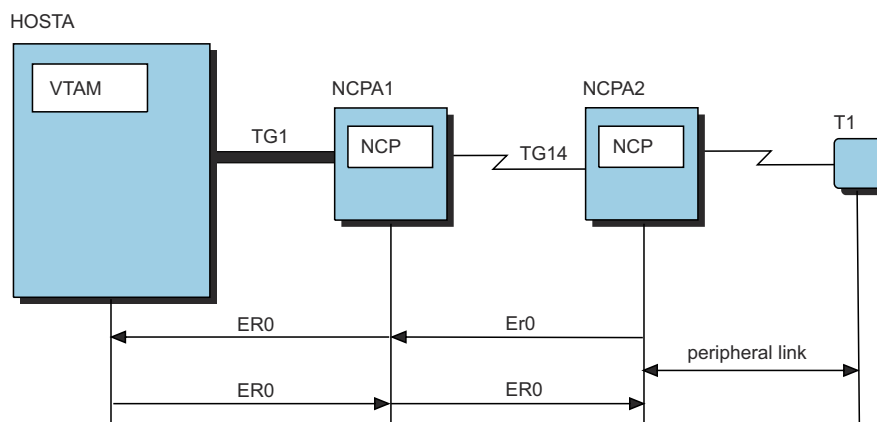


Figure 78. One explicit route in each direction

In the sample network configuration in Figure 79 on page 258, parallel SDLC links connect NCPA1 and NCPA2. Two different explicit routes between HOSTA and NCPA2 are now available in either direction. All routes take TG1 between HOSTA and NCPA1, but data flows between NCPA1 and NCPA2 over TG14 for ER0 or TG15 for ER1. The routing tables have been updated to reflect the additional explicit routes.

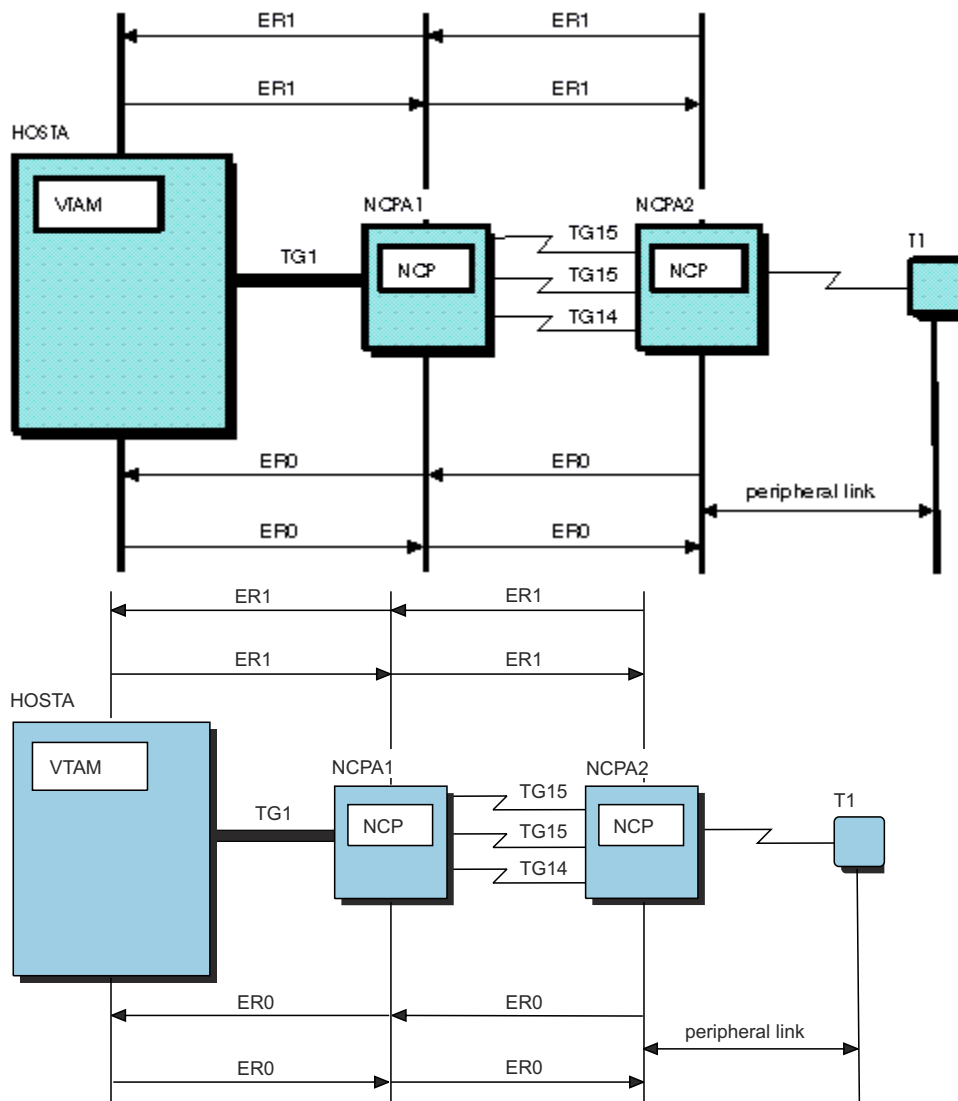


Figure 79. Two explicit routes in each direction

Logical paths

The logical path between two subarea nodes is a virtual route.

Virtual route (VR)

A virtual route is a bidirectional logical connection between two subarea nodes. At least one end of a virtual route must be in a subarea node that activates virtual routes. All hosts can activate virtual routes. NCP Version 4 Release 2 and higher releases can activate virtual routes.

Eight virtual routes numbered 0 to 7 can be defined between two subarea nodes.

One or more virtual routes must be defined for each forward-reverse explicit route pair. A virtual route places a transmission priority on data traffic using the underlying explicit routes.

You can modify the list of virtual routes and their associated transmission priorities using the virtual route selection function of the session management exit routine, or using the virtual route selection exit routine. For information about using these, see [z/OS Communications Server: SNA Customization](#).

Transmission priority (TP)

The transmission priority identifies the priority of message units flowing over an explicit route during a session. The three possible levels of transmission priority are: 0 (lowest), 1, or 2 (highest).

In general, high-priority messages are routed before low-priority messages. Within a specific transmission priority, messages are routed on a first-in-first-out (FIFO) basis.

The eight virtual routes and three levels of transmission priority provide for the possibility of 24 virtual route and transmission priority pairs between two subarea nodes. In other words, each subarea pair can have virtual routes numbered 0 through 7 between them, and each virtual route can have up to three different transmission priorities.

Route extension

A route extension is a logical connection between a subarea node and a peripheral node. A peripheral node uses local addresses for routing and requires boundary function assistance from an adjacent subarea node to communicate with a nonadjacent subarea node.

As shown in [Figure 80 on page 259](#), a virtual route must be defined for each explicit route pair. For example, between HOSTA and NCPA1, VR0 is associated with ER0 and reverse ER0, and VR1 is associated with ER1 and reverse ER1. Between HOSTA and NCPA2, VR0 is associated with ER0 and reverse ER0, and VR2 is associated with ER1 and reverse ER1. Data flows at a designated transmission priority over a virtual route. In this arrangement, data can flow over any of the virtual routes at transmission priority 0, 1, or 2.

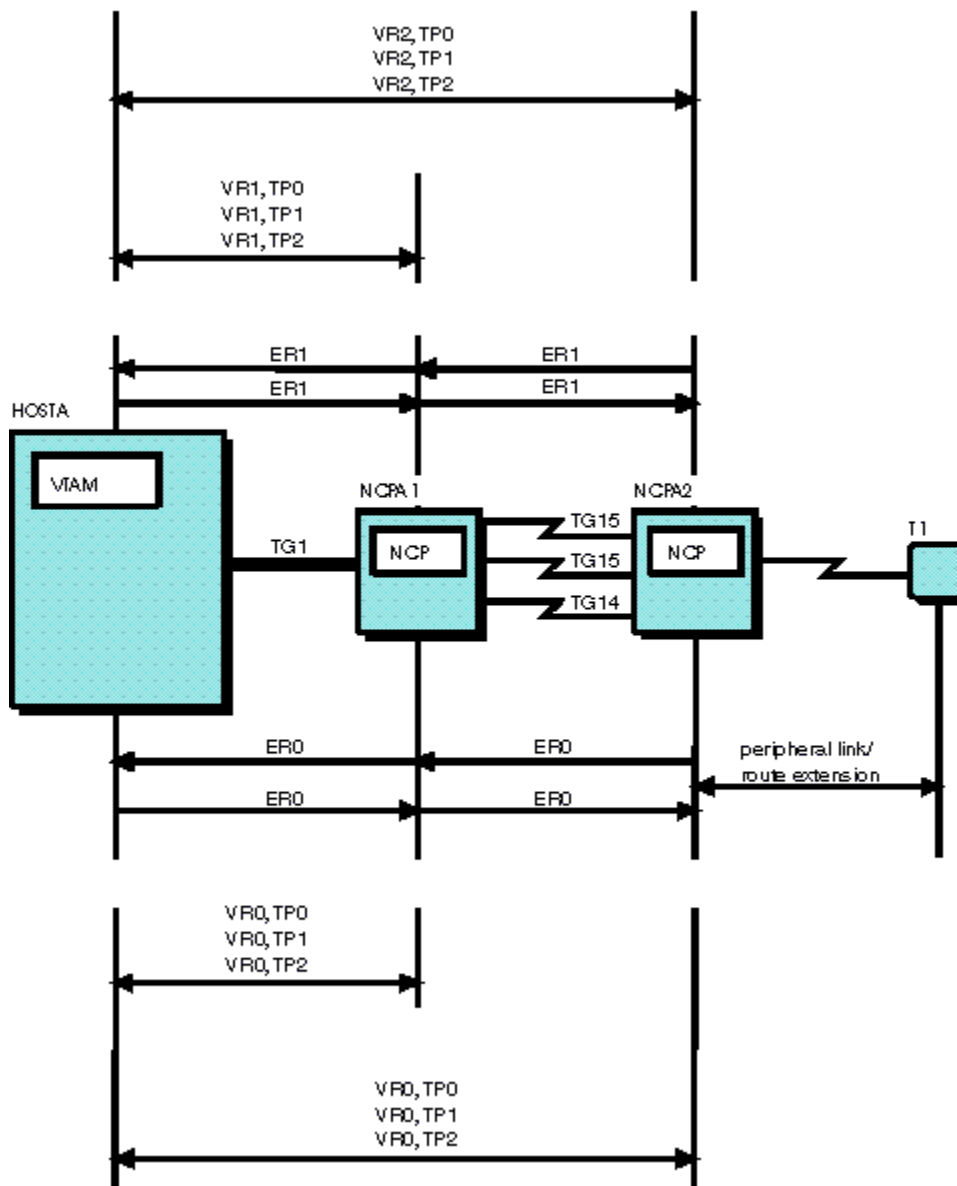


Figure 80. Virtual route and explicit route associations and transmission priority

How session traffic is assigned to a specific route

In a subarea network, data traffic for a session is assigned to a virtual route. The virtual route then maps to an explicit route over which data traffic flows at a designated priority for the session. Following are details of how session traffic is assigned to a specific route.

Logon mode table

Secondary logical units specify a logon mode entry in an LU-LU session-initiation request. Classes of service that you want to be used for an LU-LU session are specified in logon mode table entries using class of service (COS) table names similar to the following entries:

```
LOGMODE1 MODEENT COS=BATCH, ...
LOGMODE2 MODEENT COS=INTERACT, ...
:
```

If both session endpoints reside in the same subarea, the Class of Service requested by the secondary logical unit is ignored because data flowing between the two endpoints never enters the network. Otherwise, VTAM obtains the requested Class of Service from the specified MODEENT macroinstruction in the logon mode table associated with the secondary logical unit. VTAM then uses the requested Class of Service to map to the Class of Service table.

Notes:

1. If the specified logon mode entry does not exist in the logon mode table associated with the secondary logical unit, the default logon mode (ISTCOSDF) is used if it is found in the table and its use is allowed as determined by the ISTCOSDF start option.
2. If no logon mode table has been designated for a logical unit, VTAM uses the IBM-supplied logon mode table, ISTINCLM. For details on ISTINCLM, see the [z/OS Communications Server: SNA Resource Definition Reference](#).
3. If a logon mode table entry specifies a Class of Service name that is not defined in the Class of Service table, VTAM rejects the session-initiation request.

Class of service (COS) table

A class of service specifies a set of performance characteristics used in routing data between two subareas. Different classes of service can be defined in a network for different types of users and data. For example, interactive sessions can be assigned to a fast route, batch jobs to a high-bandwidth route, and sessions involving the transmission of sensitive data to a secure route. At least one class of service can also be specified for SSCP sessions.

Different classes of service are specified in a COS table. Each class of service is specified by a symbolic name (for example, INTERACT for interactive sessions or BATCH for batch sessions). The COS table lists, in order of preference, the virtual route and transmission priority pairs that are to be used for each named class of service. For example:

```
INTERACT COS VR=(1,1),(0,1),...
BATCH    COS VR=(0,0),(1,0),...
:
```

When the logon mode table entry for an LU-LU session does not contain a COS name, VTAM uses the unnamed COS entry. The unnamed COS entry is one of two special entries that you should include when you code a COS table. The other is ISTVTCOS.

Unnamed entry

An unnamed entry (the entry name consists of eight blanks) can be placed in the COS table and is used when either of the following is true:

- No class of service name is obtained from the logon mode entry for an LU-LU session.
- No ISTVTCOS entry exists in the COS table, and an SSCP session has been requested.

If the unnamed class of service entry is requested but has not been included in the COS table, VTAM uses its own Class of Service defaults.

ISTVTCOS

ISTVTCOS is the special COS entry you use to specify the routes you want to use for SSCP sessions (SSCP-SSCP, SSCP-PU, and SSCP-LU).

VTAM looks for and uses the ISTVTCOS entry any time there is an SSCP session request for another subarea. If the ISTVTCOS entry is not present, the unnamed COS entry is used.

Note: Downstream PUs and associated LUs are assigned the same VR as the boundary node for SSCP-PU and SSCP-LU sessions. For example, when an SSCP-PU session is established for a PU connected to an NCP, the VR that was chosen for the NCP SSCP-PU session will also be used for the downstream SSCP-PU session.

You need not define a COS table if the only COS names to be used are ISTVTCOS and the unnamed class of service; VTAM uses its own class of service defaults.

IBM-specified class of service defaults

VTAM has a default list of virtual route and transmission priority pairs that is used when either of the following is true:

- No COS table has been created.
- A COS table has been created, but:
 - No class of service name has been obtained from the logon mode entry associated with an LU-LU session request, and no unnamed entry has been included in the COS table.
 - An SSCP session has been requested, and no ISTVTCOS entry or unnamed entry has been found in the COS table.

The defaults consist of all 24 possible virtual route and transmission priority pairs in the following order:

```
VR0,TP0; VR1,TP0; VR2,TP0; ... VR7,TP0;  
VR0,TP1; VR1,TP1; VR2,TP1; ... VR7,TP1;  
VR0,TP2; VR1,TP2; VR2,TP2; ... VR7,TP2;
```

This default list might not be the best for your needs, and its use might not provide optimal network performance. You can replace the default list by creating a COS table with an unnamed COS entry containing the new list. This new default list is then used if no COS entry is named in an LU-LU session request or if an SSCP session is requested and no ISTVTCOS entry exists in the COS table.

Substituting class of service parameters

Usually, if VTAM receives a user session request using a COS name that is unknown to this VTAM, the session is rejected. If you want the user session to be established despite the unknown COS name, you can enable VTAM to accept the unknown COS name and use substitute parameters that are coded on another COS name. To do this, code SUBSTUT=YES on the COS name that will provide the substitute parameters.

For example, in [Figure 81 on page 262](#) assume that you are trying to set up a session between APPLA and APPLB, and APPLB is the primary LU. VTAMA specifies COSA as the COS to be used for the session. VTAMB does not have COSA coded, so VTAMB uses the parameters from COSB to set up the session because COSB has SUBSTUT=YES coded. If you display the session from either VTAMA or VTAMB, the display will indicate that COSA is being used, even though the parameters for COSB are being used. The substitution is not apparent to VTAMA. However, you can monitor the substitution in VTAMB and reject session establishments that you do not want to use the substitute parameters by coding a session management exit routine. For information about coding the exit routine, see [z/OS Communications Server: SNA Customization](#).

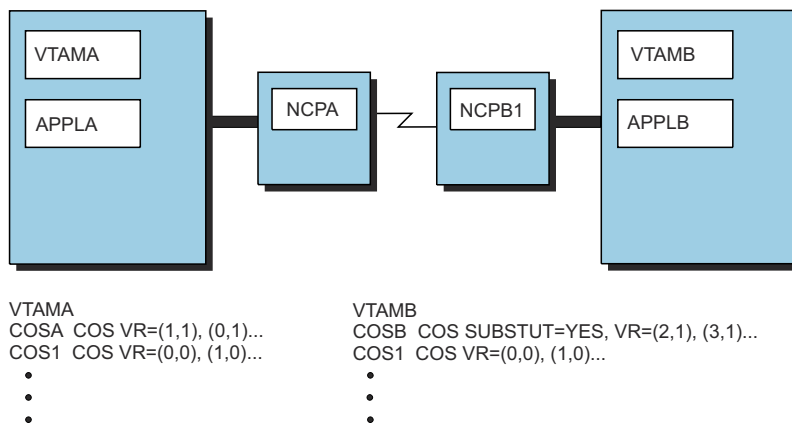
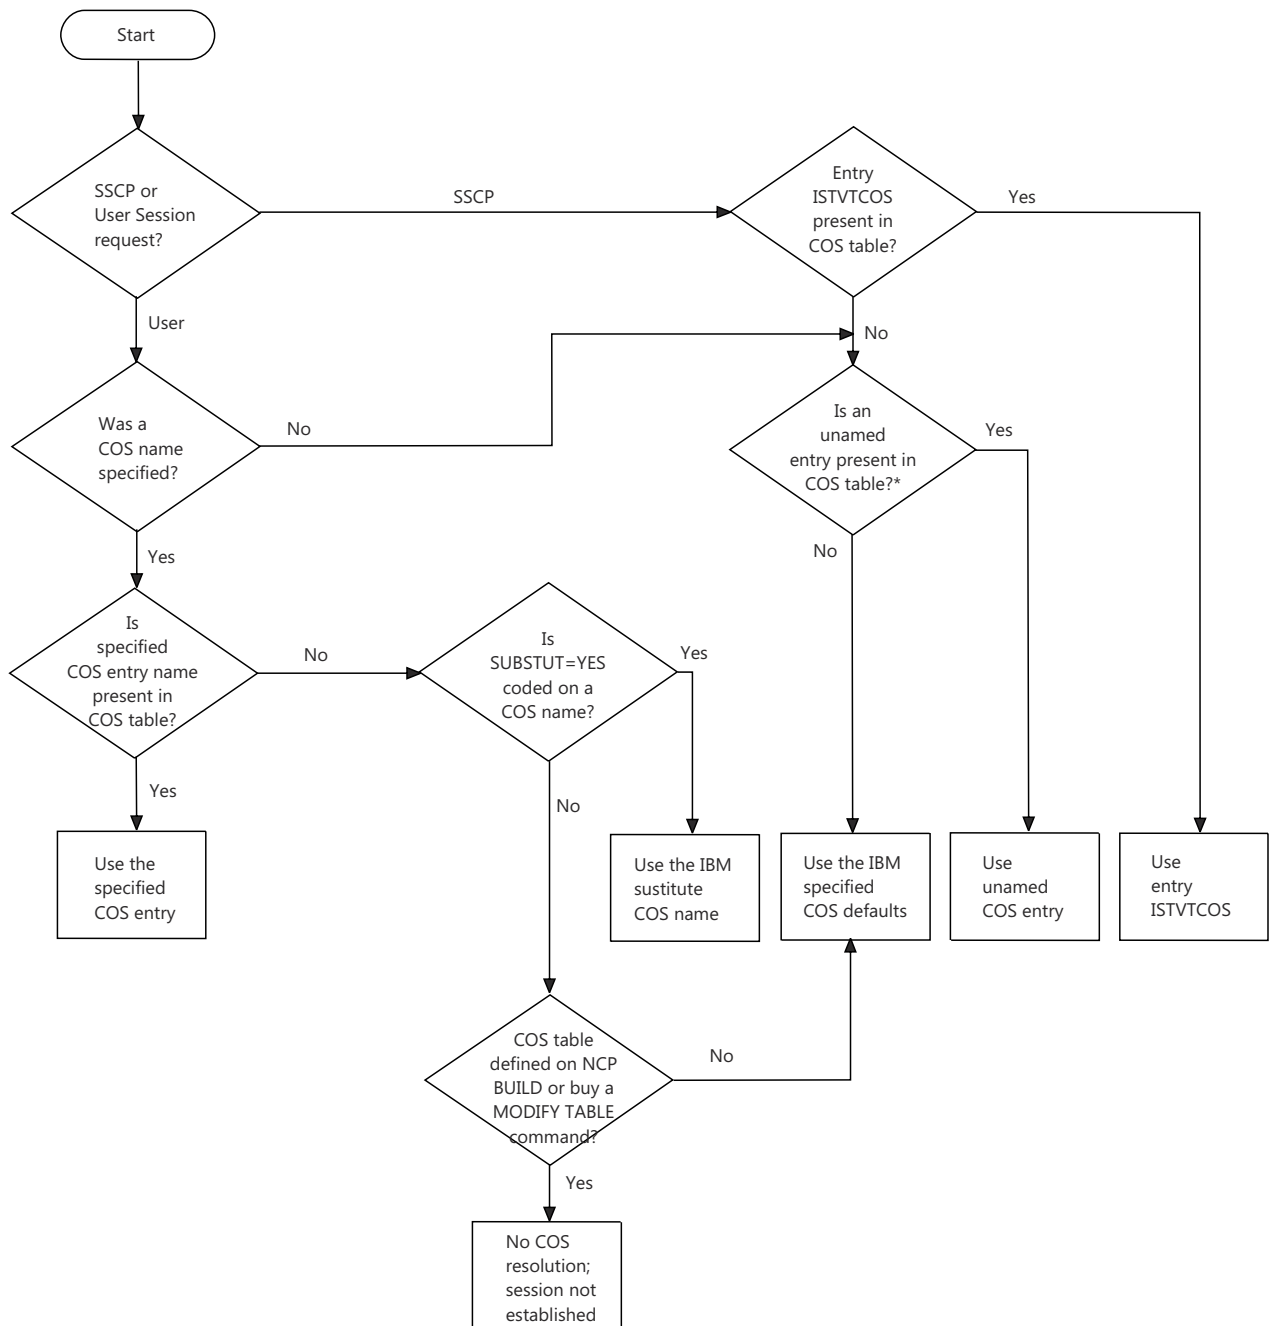


Figure 81. Class of Service substitution

Substitute class of service parameters are not passed between VTAMs. If a request is forwarded to an adjacent SSCP, the COS name that was originally requested is the name passed. The substitute COS name is not passed.

The class of service hierarchy in [Figure 82 on page 263](#) summarizes the steps VTAM takes to determine the appropriate COS table entry to use for SSCP and user sessions (or to determine to use the IBM-specified class of service defaults).

The SSCP selects the first virtual route and transmission priority pair listed for the requested class of service (or the first pair listed in the IBM-specified class of service defaults) that has been defined to the subarea of the session partner. An attempt is made to activate the virtual route.



*This is also referred to as a 'blank entry.'

Figure 82. Class of Service hierarchy

The virtual route maps to an explicit route over which data traffic flows at the designated priority for the session. If all the network resources along the explicit route are operative, the virtual route is activated and available for session use.

If the explicit route is not operative, the next virtual route and transmission priority pair for the Class of Service requested is selected, and an attempt is made to activate that virtual route. If all virtual route and transmission priority pairs for the Class of Service requested map to inoperative explicit routes, the request to establish a session is either queued or rejected. Requests for SSCP sessions are queued (if queuing is allowed), but all others are rejected.

How to plan routes in your network

Designing routes through a network ranges in complexity depending on the number of subareas, connections between subareas, and types of user sessions that your network supports. Also, decisions must be made to achieve some combination of the following objectives, which involve the following compromises:

- Maximize quantity of data transmitted
- Maximize data security
- Maximize route availability
- Minimize transmission time
- Minimize cost
- Minimize data loss or the necessity to retransmit data
- Minimize traffic congestion

NETDA/2 is a tool that can help plan for routes in your network and other network design tasks, such as:

- Designing a network
- Producing reports on designs so that you can evaluate them on performance
- Generating route statements

The following steps can also help you make the decisions necessary to plan the routes in your network:

1. Determine the Classes of Service that will exist in your network. The following questions are the types of questions you need to answer in determining the classes of service you need.
 - Will you have one Class of Service for all interactive sessions?
 - Do you want a Class of Service that provides quick response times suitable for high-priority interactive sessions and a Class of Service that provides response times suitable for low-priority interactive sessions?
 - Do you want a Class of Service that provides routes that have the best availability?
 - Do you want a Class of Service that is suitable for batch processing?
 - Do you want a Class of Service that is suitable for high-security transmissions?
2. Identify the possible explicit routes available in your network, including the possible ways parallel links between communication controller subarea nodes could be divided among transmission groups.

Dividing parallel links into multiple transmission groups creates multiple explicit routes between two nodes. Multiple explicit routes between two nodes provide the opportunity to physically separate the data traffic for different Classes of Service in your network.

Combining parallel links between communication controller subarea nodes into a single transmission group increases the probability that explicit routes using the transmission group will be operational. A transmission group that consists of parallel links will still be available even if one of the links fails, and data traffic can continue to flow on the remaining links. Only links that have similar transmission characteristics (for example, high-transmission speed) are normally placed in the same transmission group.

3. Analyze the physical characteristics of each identifiable explicit route, such as:
 - Length
 - Transmission time
 - Lowest link capacity in the route
 - Lowest transmission group capacity in the route
 - Lowest security level in the route
 - Operational probability

4. Match the Classes of Service you selected to explicit routes having appropriate characteristics. Remember to consider the characteristics of a route between each individual subarea pair. For example, for a Class of Service requiring a fast route, choose the fastest route between each subarea pair along the route.

Also, be sure to consider alternate routes for each Class of Service. If a subarea node or transmission group fails, all routes passing through that node or transmission group become inoperative and existing sessions are disrupted. An alternate route is a route that is available if the route being used becomes inoperable. If you want to ensure that you always have a route available, you should choose alternate routes that do not pass through the same subarea nodes or transmission groups as the routes that they back up.

5. Identify a virtual route for each selected explicit route.
6. Create an ordered set of virtual route and transmission priority pairs for each Class of Service.

In creating this ordered set, you assign a transmission priority to a virtual route which, together with the physical characteristics of the explicit route, determines the suitability of a particular virtual route for a designated Class of Service. Remember that multiple sessions can use the same virtual route (with the same or a different Class of Service).

For example, for a Class of Service requiring faster data transmission and predictable response times (a Class of Service suitable for interactive sessions), you can list virtual routes assigned to explicit routes that have fewer physical elements before listing virtual routes assigned to longer explicit routes. The shorter routes would be used first, and the longer routes would be used only for backup purposes.

However, you might have interactive users with differing Class of Service requirements. The virtual routes for each of these Classes of Service could be identical to those above, and the Class of Service would be realized through the transmission priority. For the interactive users requiring faster data transmission and predictable response times (for example, an inquiry-response Class of Service), you would list virtual routes having a higher transmission priority than the virtual routes for interactive users where slower data flow is acceptable (for example, a data-collection Class of Service). In other words, sessions for different kinds of applications can be in progress over a given explicit route, and transmission priorities determine the order of transmission.

Suppose you have two Classes of Service for low priority and higher priority batch sessions. The higher-priority batch sessions would be assigned to virtual routes with a higher transmission priority than the virtual routes for the lower-priority batch sessions. For both batch Classes of Service, you would probably list high-bandwidth explicit routes before listing the explicit routes used for interactive sessions. The shorter routes used for interactive sessions could be used as alternate routes for batch sessions. The batch sessions would likely be assigned a transmission priority lower than the transmission priority of the interactive sessions using the same virtual route.

Note: In migration environments, you should include VR 0, TP 0 in each COS table entry as an identifier for default routes containing nonextended addressing nodes.

7. For details on coding a COS table to hold the ordered virtual route and transmission priority pairs for each class of service, see the [z/OS Communications Server: SNA Resource Definition Reference](#).
8. As described in [“How session traffic is assigned to a specific route” on page 260](#), VTAM processes the virtual route and transmission priority pairs in a COS table entry sequentially, starting with the first entry. If you require a more precise selection algorithm, use the session management exit routine or virtual route selection exit routine to change the virtual route selection process for LU-LU sessions.

For example, if you provide a virtual route selection exit routine, your routine receives the ordered virtual route list as a parameter from VTAM. Your routine can modify the list. VTAM then uses your reordered list of virtual routes to select a virtual route for the session.

If you have provided a virtual route selection exit routine, VTAM calls it whenever a session between a primary logical unit in the VTAM subarea and a logical unit in another subarea is about to be established. The virtual route selection exit routine is not called if both logical units are in the same VTAM subarea.

For more information about coding a virtual route selection exit routine or using the session management exit routine to change the virtual route selection process for LU-LU sessions, see [z/OS Communications Server: SNA Customization](#).

9. For each subarea node, use PATH definition statements to create the path definitions necessary for defining the explicit routes and virtual routes that exist between the subarea node and other subareas. For example, the PATH statements needed in HOSTA of [Figure 80 on page 259](#) are:

```
PATH1  PATH  DESTSA=NCPA1,
              ER0=(NCPA1,1),
              ER1=(NCPA1,1),
              VR0=0,
              VR1=1
PATH2  PATH  DESTSA=NCPA2,
              ER0=(NCPA1,1),
              ER1=(NCPA1,1),
              VR0=0,
              VR2=1
```

From path definitions like these, VTAM builds a routing table for a subarea node similar to the table for HOSTA shown in [Figure 79 on page 258](#). In creating your path definitions, list virtual routes only in those subareas that are going to activate them. For NCP subareas that cannot activate virtual routes, the virtual routes ending in those subareas are dynamically defined to those subareas as part of route activation.

Coding paths for a multiple-domain environment is similar to coding paths in a single-domain environment. Communication between domains requires paths to be defined between VTAM subareas. The PATH definition statements that are used to define routes between VTAM and NCP subareas in a network are also used to define routes between the VTAM domains. For each subarea, the PATH definition statements define the explicit and virtual routes that connect the various VTAM domains.

VTAM can function as an intermediate routing node. VTAM receives the data from the adjacent subarea and transmits it to the next subarea on the network route as defined by the PATH definition statements. The IRNSTRGE start option is used to define the maximum size of the virtual area in VTAM storage that can save host intermediate routing node (IRN) transmissions. If IRNSTRGE=0 (the default) is used, the amount of storage is unlimited.

For further details on PATH statements or the IRNSTRGE start option, see [z/OS Communications Server: SNA Resource Definition Reference](#).

10. One or more path definition sets must be activated for VTAM to know how to route cross-subarea session traffic. A path definition set is activated with a VARY ACT command. When a file containing a set of PATH definition statements is activated, it defines explicit routes and virtual routes to VTAM. At that point, the routes are defined but not yet active.

You can activate additional path definition sets to add or modify route definitions when necessary. You can define new routes and redefine or delete inactive routes. Existing route definitions not affected by a given path definition set activation are unmodified and remain in effect. VTAM rejects any attempt to modify the definition of any previously defined route that is not in the inactive state.

A path definition set is not a major node and cannot be deactivated.

11. In addition to the activation of path definitions, all physical elements included in an explicit route (subarea nodes and transmission groups) must be activated before an explicit route is ready for activation.

An explicit route is operational when physical connectivity between the endpoints is established. Path information is exchanged between each subarea when the connection between the subareas is activated. A virtual route becomes available for activation after its corresponding explicit route is active.

VTAM automatically attempts to activate explicit and virtual routes when they are first needed for a session. A virtual route remains active until all sessions assigned to it have ended, at which time it is automatically deactivated. The deactivation of a virtual route has no effect on the status of

its associated explicit route. When active, an explicit route is never deactivated, but it can become inoperative because of hardware failure or deactivation of physical elements within the route.

12. You can use the DISPLAY ROUTE command to display information about the explicit routes and virtual routes between the host subarea and a given destination subarea. You can also use this command to test whether one or more explicit routes to a given destination are operative (see [“Displaying and testing routes”](#) on page 488).
13. Remember that the requested Class of Service for a session is obtained from the logon mode table associated with the secondary LU. For details on creating a logon mode table and associating LUs with your table, see [z/OS Communications Server: SNA Resource Definition Reference](#). If you do not create a logon mode table and also make the association between an LU and that table, VTAM obtains the LU requested Class of Service from the IBM-supplied logon mode table, ISTINCLM.

How VTAM handles network and subarea addressing

For VTAM to control routing in a subarea network or within a composite network node, VTAM must know the location of its resources, such as logical units, physical units, and other SSCPs. VTAM uses element addresses, in conjunction with the subarea address, to identify the location of resources (also known as network addressable units). The subarea address indicates in which subarea the resource is located; the element address indicates the unique address within the subarea.

Minor nodes, such as an application program or logical unit, require at least one element address. Some minor nodes, such as local non-SNA devices and application programs that use parallel sessions, require more than one element address. This can increase the number of element addresses that is used in a subarea.

The different types of addressing structures are:

Preextended network addressing

Used by releases before VTAM Version 3 Release 1. Enables you to define 255 subareas and 254 elements.

Extended network addressing (ENA)

Used by VTAM Version 3 Release 1 and subsequent releases before VTAM Version 3 Release 2. ENA supports addressing to 255 subareas and extends element addressing to 32 768 elements. Beginning with VTAM Version 4 Release 1, host element addressing was extended to 65 535 elements.

Extended subarea addressing

Used by VTAM Version 3 Release 2 (with compatibility PTF) and subsequent releases. Extended subarea addressing increased the size of the subarea address to 65 535 and the number of explicit routes for each destination subarea to 16.

Note: Applications are assigned a low-order primary LU address upon activation. This low-order address may be used in cases where a high-order address could have been used.

Compatibility between nonextended and extended network addressing nodes

You can use the MAXSUBA start option or the MAXSUBA operand in NCP to enable subareas with different addressing structures to communicate. For a network containing nonextended addressing nodes, the MAXSUBA start option specifies the highest subarea value that can be used throughout the network. A MAXSUBA of 63, for example, defines a network with up to 63 subareas and 1024 elements in each subarea. Code the MAXSUBA start option in the start option list if you want VTAM to communicate with nonextended addressing nodes, or if the MAXSUBA operand is coded in an NCP that VTAM will communicate with.

Extended network addressing is a network addressing structure that increases the size of the subarea address up to 255 in conjunction with up to 32 768 elements. VTAM extends host element addressing to 65 535 elements.

Nonextended addressing nodes can continue to use this structure. They can, however, participate in an extended network addressing network only when the following two conditions are met:

- A compatibility program temporary fix (PTF) must be installed on all nonextended addressing nodes that communicate with extended network addressing nodes.
- All nodes that communicate with nonextended addressing nodes must define a MAXSUBA. Extended network addressing nodes need this information to decode 16-bit network addresses. The value specified for MAXSUBA must be the same at all nodes.

In addition, the following restrictions apply when extended and nonextended addressing nodes coexist in the same network:

- VTAMs that are nonextended addressing nodes cannot communicate with any logical unit whose element address exceeds the maximum element number defined by the subarea/element split.
- VTAMs that are nonextended addressing nodes cannot communicate with any node whose subarea is greater than the MAXSUBA.

Figure 83 on page 268 illustrates an element address incompatibility. Domain 1 contains nodes that use extended network addressing. Domain 2 contains nodes that use nonextended addressing node addresses. All nodes have defined a MAXSUBA of 63. Although all of the subareas are within the MAXSUBA, some of the element addresses exceed the maximum, which is 1023. Following are two address constraints on this network:

- In Domain 1, the terminals T1 and T3 can communicate with any subarea because their element addresses (100) are less than 1024. T2 and T4, however, cannot communicate with host A40M because their element addresses (1500 and 2000) are greater than 1023.
- In Domain 2, T5 and T6 can communicate with either host. However, they cannot communicate with any application program in host A60M whose element address is greater than 1023.

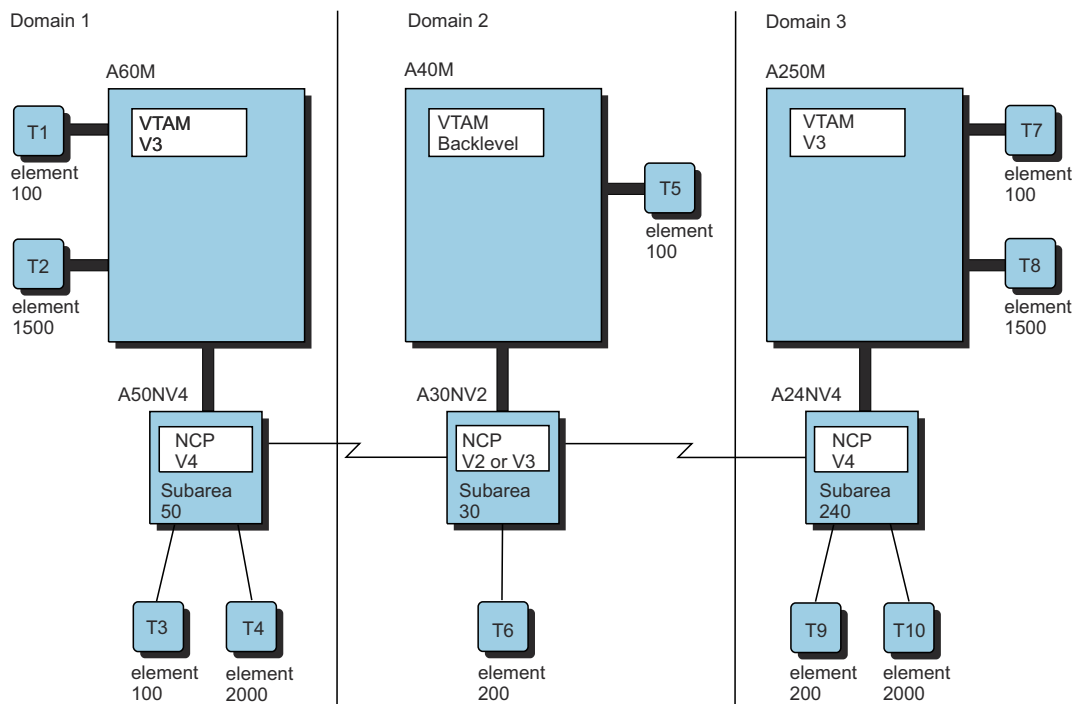


Figure 83. Element and subarea address incompatibility in multiple-domain environment

These limitations can be overcome through SNA network interconnection (SNI), which can be used to place nonextended addressing nodes and extended network addressing nodes into separate networks. See Chapter 18, “Connecting multiple subarea networks,” on page 439 for information about implementing SNI.

Although it is not recommended, a network containing nonextended addressing nodes can have subareas whose subarea number is greater than the MAXSUBA. The nonextended addressing nodes are unable to communicate with these subareas. In this situation, the following constraints apply:

- A nonextended addressing node VTAM cannot appear in any route whose endpoints have a subarea address greater than MAXSUBA.
- A nonextended addressing node VTAM cannot be adjacent to a node whose subarea is greater than MAXSUBA.
- A nonextended addressing node NCP can be adjacent to a node whose subarea is greater than MAXSUBA, but only if the NCP is used as an intermediate routing node (IRN). This requires a PTF.

Table 37 on page 269 and Table 38 on page 269 show the adjacent and endpoint node subarea requirements.

<i>Table 37. Adjacent node subarea requirements for multiple-domain environment</i>			
	Pre-ENA VTAM	ENA VTAM subarea ≤ MAXSUBA	ENA VTAM subarea > MAXSUBA
Pre-ENA NCP	Yes	Yes	Yes ³
ENA VTAM Subarea < MAXSUBA	Yes	Yes	Yes
ENA VTAM Subarea > MAXSUBA	No	Yes	Yes
Notes: <ol style="list-style-type: none"> 1. Yes: These two nodes can be endpoints. 2. No: These two nodes cannot be endpoints. 3. Pre-ENA is used as an IRN only. 			

<i>Table 38. Endpoint node subarea requirements for multiple-domain environment</i>			
	Pre-ENA VTAM	ENA VTAM subarea ≤ MAXSUBA	ENA VTAM subarea > MAXSUBA
Pre-ENA NCP	Yes	Yes	No
ENA VTAM Subarea < MAXSUBA	Yes	Yes	Yes
ENA VTAM Subarea > MAXSUBA	No	Yes	Yes
Notes: <ol style="list-style-type: none"> 1. Yes: These two nodes can be endpoints. 2. No: These two nodes cannot be endpoints. 			

Figure 83 on page 268 illustrates these constraints. The third domain contains extended network addressing nodes whose subarea addresses are greater than MAXSUBA.

The following communication restrictions apply to the network in Figure 83 on page 268:

- None of the Domain 2 terminals can have a session with host A250M. Conversely, the Domain 3 terminals cannot have a session with host A40M.
- All of the terminals in Domain 3 can communicate with any application program in Domain 1, and all of the terminals in Domain 1 can communicate with any application program in Domain 3.
- NCP A30NV2 is a nonextended addressing node, but can still act as an IRN between Domains 1 and 3. However, the NCP terminal, T6, cannot communicate with Domain 3.

As with the element address problem, subarea address incompatibilities can be avoided by using SNI to place nodes in separate networks. See Chapter 18, “Connecting multiple subarea networks,” on page 439 for information about implementing SNI.

Compatibility between extended network and extended subarea addressing nodes

Whereas extended network addressing supports a maximum subarea address of 255, extended subarea addressing allows you to specify subarea addresses higher than 255. Subarea addresses can be as high as 65 535.

However, defining networks with large numbers of subareas can increase the storage requirements for VTAM and NCP. To control the maximum number of subareas in a network, use the MXSUBNUM start option and the NCP operand SALIMIT to specify the maximum subarea size of the network. You can define VTAMs and NCPs in the network with varying values for the maximum subarea limit. However, the actual subarea addresses used must be within a common address range for the subarea nodes to communicate. Processors and communication controllers with levels of VTAM and NCP that do not support extended subarea addressing can exist in the same network with other processors and communication controllers that do. For subarea node communication to be supported, the actual subarea addresses used must be within 255 or the MAXSUBA value, whichever is smaller.

Migration can be simplified by installing levels of software that support extended addressing on all subareas in the network, with the maximum subarea limit specified to accommodate planned growth, before adding any subareas with a higher subarea address value. Failure to do so results in subareas that no longer have connectivity to other parts of the network.

The maximum subarea limit should, in general, be constant throughout the network. If the subarea limit cannot be constant, you can try different values with no loss of connectivity until a subarea address is used that is higher than the lowest subarea limit in the network. When moving to a lower subarea limit value, you should lower all subarea numbers before changing the subarea limit value.

The following subarea addressing restrictions apply to a multiple-domain network unless you use SNA network interconnection (SNI) to circumvent the addressing constraints:

- Adjacent subareas can support different maximum subarea values and still communicate provided their actual subarea addresses fall within their common range. For example, a nonextended subarea addressing VTAM can communicate with an extended subarea addressing NCP that supports an SALIMIT value greater than 255 if the NCP subarea address is no greater than 255.
- A subarea whose address is above the addressing supported by other adjacent subareas on a route cannot communicate to or through those subareas.
- Subareas can communicate through other subareas with an address higher than they support provided there is an intermediate subarea with an address acceptable to both. For example, a nonextended subarea addressing VTAM could communicate through an NCP with a subarea address of 300. There would, however, have to be an intermediate VTAM or NCP supporting extended subarea addressing with an address no greater than 255 and an SALIMIT of at least 300.
- An SSCP cannot activate any subareas with addresses greater than its addressing capabilities.
- For an extended subarea addressing NCP to successfully send a CONTACT request unit to an adjacent extended subarea addressing NCP with an address above 255, the link station must be owned by an extended subarea addressing VTAM.

With extended subarea addressing, VTAM and NCP support the use of 16 explicit routes. You can define up to 16 explicit routes for each destination subarea. However, for an explicit route to have a number greater than 7, all subareas on the route must support extended subarea addressing, and all NCPs must have ERLIMIT=16 coded on the BUILD definition statement.

Virtual route pacing

The term virtual route pacing applies to the monitoring and flow control mechanisms that are used for virtual routes. On a network-wide basis, subarea nodes (VTAM and NCP) continuously monitor the amount

of traffic in the network. If congestion occurs, these subarea nodes can limit the amount of data they send over the affected virtual routes until the congestion clears. The monitoring of data flow and limiting of data occurs automatically, requiring no action by users or domain operators.

Activation of a virtual route includes initializing the pacing count. This count, the virtual route window size, indicates the number of path information units (PIUs) that are allowed to flow on a virtual route before the subarea node receiving the PIUs authorizes the sending of more data. The set of PIUs that can be sent is called a window. This value fluctuates between a minimum and a maximum value depending on network congestion. (Network congestion is determined by how many messages are queued.) You can affect how VTAM does virtual route pacing through the PATH definition statement or the VR pacing window size calculation module. The minimum and maximum values can be specified on the NCP PATH definition statements and the VTAM PATH definition statements that define the virtual routes.

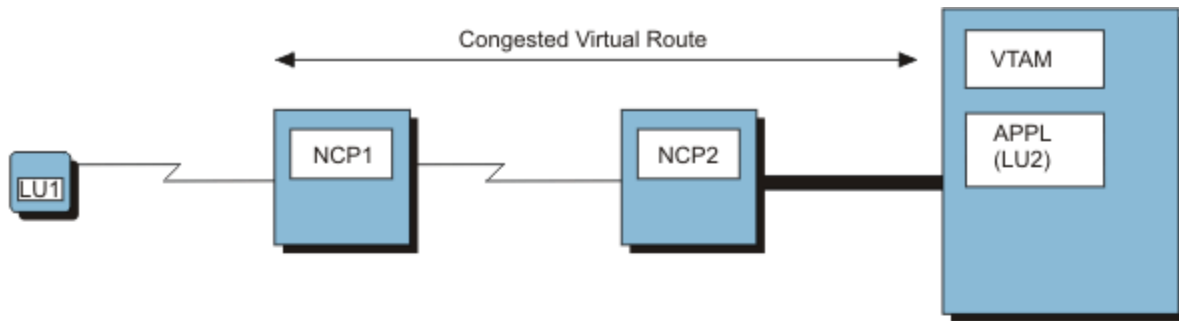
Degrees of congestion

The network detects two degrees of congestion: minor and severe. Minor congestion is detected when the amount of data queued for a specific transmission priority to a particular transmission group exceeds a user-specified value. Severe congestion is detected when the aggregate amount of data for all transmission priorities exceeds the total user-specified value for a transmission group. At the first, less severe level, the number of PIUs allowed to be sent in a window is decreased by one until congestion clears. At the severe level, the window size is immediately decreased to its minimum value. These two levels can be specified with NCP transmission group flow control threshold parameters. These threshold parameters can be coded on the ER operand on the NCP PATH definition statement.

VTAM detects only severe congestion. The reset window indicator is set in the PIU if VTAM cannot immediately get an IO buffer, either because of a shortage or because XPANPT is not high enough to allow expansion to complete before all the remaining buffers are allocated.

Local flow control and severe congestion

Virtual route pacing is supplemented by a local flow control function that limits the amount of data that users within a subarea can present to the network for transmission. [Figure 84 on page 272](#) illustrates local flow control.






<p>Example A. Local Flow Control in NCP1 for Sessions that use Pacing:</p> <div style="display: flex; justify-content: space-around; align-items: flex-start;"> <div style="text-align: center;"> <u>LU1</u> MSG 1  MSG 2  . . . MSG n  n=LU1's pacing count. </div> <div style="text-align: center;"> <u>NCP1</u> NCP1 waits to send a session pacing response until the virtual route is no longer congested. </div> </div>		<p>Example B. Local Flow Control in VTAM:</p> <div style="display: flex; justify-content: space-around; align-items: flex-start;"> <div style="text-align: center;"> <u>VTAM</u> VTAM waits to complete the SEND request until the virtual route is no longer congested. </div> <div style="text-align: center;"> <u>LU2</u> . . . SEND </div> </div>	
<p>Example C. Local Flow Control in NCP1 for Sessions that do not use Pacing:</p> <p>NCP1 waits to read from the device (LU1) until the virtual route is no longer congested.</p>			

Figure 84. Local flow control

Suppose, for example, that the virtual route between NCP1 and VTAM is congested and that VTAM has been notified of the congestion. If NCP1 does not give VTAM permission to send another window of data, the virtual route is held. VTAM then keeps the windows of data in IOBUF buffer space until the data can be passed.

If LU1 is an SNA device and session pacing is being used on the LU-LU session, NCP1 withholds the session pacing response going to LU1 until the virtual route is no longer congested. (See Example A in Figure 84 on page 272.) Because virtual route pacing controls only normal-flow requests, LU1 can still send responses to previously received requests and expedited flow requests. Because those RUs can also congest the virtual route, NCP1 holds them until the congestion clears.

If LU1 is a non-SNA device (or an SNA device that is not using session pacing on the LU-LU session), NCP1 does not read any input from the device until the virtual route is no longer congested. (See Example C in Figure 84 on page 272.)

When the application program (LU2) in Host1 sends data to a device attached to NCP1, VTAM holds the data in IOBUF buffer pool space until it is informed that the virtual route is no longer congested. After the virtual route is blocked and VTAM holds the windows of data, VTAM queues the application program SEND request and does not read the data from the user buffer. (See Example B in Figure 84 on page 272.) The

direction of data flow is paced independently. One direction might be congested while the other direction is not. [Figure 84 on page 272](#) shows a virtual route that is congested in both directions.

Note: If VTAM determines that sessions with one of its local devices are contributing to congestion over a route, it shuts off data transmission from the device. Therefore, if there are a number of SNA devices attached to a channel-attached cluster controller, sessions for all the devices are shut down, regardless of whether their sessions are being conducted over the congested route or another route. To avoid this, you might want to specify inbound session pacing for all channel-attached SNA devices so that they are each paced separately. See [“Session-level pacing” on page 217](#).

Window sizing

The IBM-supplied algorithm for window size calculation is designed to work with the route pacing algorithm used in the network. It is appropriate for most installations and configurations. VTAM calculates the minimum and maximum sizes of virtual route pacing windows based on the link protocol and the explicit route length (that is, the number of transmission groups in the explicit route used by the virtual route). While the virtual route is being used to transmit data, adjacent subarea nodes on the route automatically adjust the window sizes within the minimum and maximum limits according to traffic conditions along the route.

However, after tuning VTAM and analyzing traffic patterns and resource capabilities, you might want to choose your own bounds or code a replacement routine that sets the window sizes to different values than the ones supplied by IBM. This exit routine is appropriate for systems where the number of resources could vary considerably from one day to the next. For example, you might find one or more resources whose capacities are not consistently used. This might warrant increasing the window sizes. Decreasing the window sizes is less likely to be useful because network flow control protocols are designed to prevent congestion, and setting window sizes too small could reduce traffic flow considerably.

If you use the Information Management System (IMS), you might want to code a replacement virtual route-pacing window-size calculation exit routine to reduce the number of virtual route pacing responses that VTAM processes for every IMS transaction. For more information about coding the exit routine for IMS, see [z/OS Communications Server: SNA Customization](#).

You can also specify default minimum and maximum window sizes on the `VRPWSnn` operand of the `PATH` definition statement when defining a virtual route. This is the simplest way to choose the default and is sufficient for most needs. If you do not specify `VRPWSnn`, VTAM uses the default algorithm described previously.

The virtual route-pacing window-size calculation exit routine can be used in addition to `VRPWSnn` operands (coded on the VTAM and the NCP `PATH` definition statements) to alter the VTAM virtual route pacing window sizes. This exit routine is used to specify the bounds for virtual route pacing windows. A virtual route pacing window represents the quantity of path information units (PIUs) that can be transmitted on a virtual route before a virtual route pacing response is received. This response indicates that the virtual route receiver is ready to receive more PIUs on the route. The exit routine is called when a virtual route is activated. It returns the minimum and maximum values for the window of the virtual route.

Parallel sessions using parallel transmission groups

If the application program supports parallel sessions, each session can use a different class-of-service entry and can map to a route that uses a different channel path.

For example, in [Figure 85 on page 274](#), because the type 2.1 peripheral node is connected to an NCP with parallel channel-link transmission groups (TG1 and TG2), the independent logical unit (APPL21) can establish parallel sessions with CICS in the host, and each session can use a different route. However, each session must use a different Class of Service entry (such as COS1 and COS2). These different Class of Service entries map to different virtual routes (such as VR1 and VR2) using the same or different transmission priorities (in this case, the same, TP1). As shown in [Figure 85 on page 274](#), the different virtual routes map to the explicit routes for each session.




Figure 85. Parallel sessions using parallel transmission groups

Dynamic path update

Dynamic path update enables the VTAM operator to add NCP path definitions dynamically without regenerating and reloading the NCP. Using the same operator facilities, you can also dynamically add VTAM path definitions. The explicit route definitions in VTAM and NCP can be replaced or deleted only if the explicit route is inoperative (a status of INOP).

Note: If you are not attempting to modify an ER mapping to a TG, the ER state is not checked and the ER can be either operative or inoperative.

Dynamic path update can also be used to change inactive (a status of INACT) or undefined VRs, path definition values for virtual route pacing window size on inoperative routes, and transmission group thresholds on operative routes.

To use dynamic path update, at least one route must be defined and operative to establish the necessary SSCP-PU session between VTAM and each NCP. The dynamic path update changes are transmitted to NCP on this session. Dynamic path update members are filed in the VTAM definition library. A dynamic path update member can contain path specifications that are added, replaced, or deleted from one or more NCP or VTAM subareas. A set of path update specifications that is targeted to a specific subarea node within a dynamic path update member is called a dynamic path specification set. A dynamic path update member can contain one or more dynamic path specification sets.

The dynamic path update member in the following sample can be used to add NCP3 to the network shown in Figure 86 on page 275.

```
* Dynamic Path Specification Set for VTAM 1 (subarea 1)
VTAM1  VPATH NETID=NETA
      PATH  DESTSA=3,
          ER0=(2,1),
          VR0=0
* Dynamic Path Specification Set for NCP 2 (subarea 2)
NCP2    NCPPATH NETID=NETA
      PATH  DESTSA=3,
          ER0=(3,10)
```

Subarea 1




Figure 86. Sample single-domain network

Note that the name on the NCPPATH definition statement identifies the NCP for which the PATH statements are to be processed.

A dynamic path update can be triggered by any of the following methods:

- Activating a dynamic path update member using the VTAM configuration list or the VARY operator command. An active SSCP-PU session must exist between VTAM and the NCP whose routes are being dynamically updated, and the explicit routes being changed must be inoperative.
- Specifying the NEWPATH operand when activating an NCP with an operator command. This causes VTAM to activate the dynamic path update member immediately following NCP activation.
- Activating an NCP that has the NEWPATH operand specified on the PCCU definition statement for this host. This causes VTAM to activate the dynamic path update member immediately following NCP activation.

Note: Because dynamic path update changes do not alter VTAM or the NCP load module, you should use a method that ensures that these changes are reinstated if VTAM is restarted or if NCP is reloaded.

The dynamic path update member in the preceding example could be filed in the VTAM1 definition library. When the operator activates the dynamic path update member, VTAM1 identifies the dynamic path update member because the name in the VPATH definition statement matches its SSCPNAME (VTAM1). VTAM1 transmits the dynamic path update member to NCP2 using the active session path (SSCP-PU session).

Note: You should file the dynamic path update member for each subarea in separate dynamic path update members to avoid possible performance deterioration.

The following restrictions are placed on coding dynamic path update members and dynamic path specification sets:

- A dynamic path update member can include path definitions for only one VTAM host. If an incorrect definition statement or VPATH or NCPPATH definition statement is encountered, VTAM discards all PATH definition statements that follow the definition statement that is not valid up to the next recognizable VPATH or NCPPATH definition statement and issues a warning message. PATH definition statements that were already processed are not affected.
- When a dynamic path update member is activated using the VTAM configuration list or the VARY operator command, an NCPPATH definition statement is considered to be not valid if the NCP name on the NCPPATH definition statement does not match the resource name of an active NCP.
- When a dynamic path update member is activated by specifying the NEWPATH operand, an NCPPATH definition statement is considered to be not valid if the NCP name on the NCPPATH definition statement does not match the resource name of the NCP being activated.
- A dynamic path update member can include the DELETER operand if it applies to the host activating the dynamic path update member.

VTAM saves the name of the dynamic path update member in the NODELST data set (the VTAM configuration restart data set). The actual dynamic path specification sets are not saved in a VSAM configuration restart file.

Note: Extreme caution must be used when deleting or replacing a route for a node that is an intermediate routing node for other routes.

The NCPPATH definition statement identifies the NCP and the network to which the dynamic path update member applies in a dynamic path update member.

Chapter 13. Application programs

Each application program must be defined within an application program major node. Each application program represents a minor node.

You define an application program major node with one VBUILD definition statement, and you define application program minor nodes with one APPL definition statement for each application program in the major node. A sample of definition statements for an application program major node is provided. Each APPL definition statement represents a single logical unit to VTAM.

```
APPLNODE VBUILD TYPE=APPL
A50PAYRL APPL    AUTH=ACQ,          PERMIT APPLICATION TO ACQUIRE LUS
                                CONCURRENT APPLICATION SESSIONS
                                LOGON MODE TABLE=APPLICATION IS SLU
                                ENABLE LU 6.2 SUPPORT
                                ENABLE PARALLEL SESSION FOR LU 6.2 SUPPORT
                                NETWORK LU NAME=A50PAYRL
                                VTAM APPLICATION APPLID=PAYROLL
:
A50ACCTS APPL    ATNLOSS=ALL,       ENABLE ATTN EXIT FOR ALL LU 6.2 SESSIONS
                                ENABLE LU 6.2 SUPPORT
                                ENABLE PARALLEL SESSION FOR LU 6.2 SUPPORT
                                ACCEPT ALREADY VERIFIED LU 6.2 CONVERSATIONS
                                NETWORK LU NAME=A50ACCTS
                                VTAM APPLICATION APPLID=ACCOUNTS
:
```

You can also code a model application program definition, which can be used as the definition for one or more VTAM application programs. For some application programs, such as Telnet or TSO/VTAM, using model application program definitions can save on resources (such as host network addresses) and is recommended.

Naming an application program

You can specify two different names for an application program, one required and one optional. The required application program name is the one that you specify in the label field of the APPL definition statement. The optional application program name is the one that you specify on the optional ACBNAME operand of the APPL definition statement.

You can use wildcard characters in the name you specify in the label field and on the ACBNAME operand. Using wildcard characters enables you to define a model application program definition, which can be used as the definition for one or more VTAM application programs. For more information about defining model application program definitions, see [“Model application program definitions” on page 278](#).

If an application program name is not specified on the APPLID operand of the ACB macroinstruction, VTAM uses a name supplied by your operating system. If the name supplied by your operating system does not match a name on the APPL definition statement, the application program cannot open its access method control block (ACB) and is therefore not identified to VTAM.

If the application program is started by a job step that is a procedure invocation, the name used by the operating system is the procedure step name in the job control language (JCL) for the application program. Otherwise, it is the job step name in the JCL.

Note: TSO/VTAM and logon manager are VTAM application programs and must be defined to VTAM. For the specific coding requirements for the APPL definition statement that defines TSO/VTAM, see [Appendix A, “TSO/VTAM,” on page 557](#). For the specific coding requirements for the APPL definition statement that defines logon manager, see [Appendix D, “Logon manager,” on page 585](#).

If you are using an application program that needs information about the environment in which it is running, you can code macroinstructions in that application program to obtain information about:

- VTAM version and release levels

- VTAM component and feature identifiers
- Application program functions provided by your release of VTAM
- Application program network name
- Application program ACB name
- HOSTPU name
- Network ID for your network

For information about application program pacing, see [“Application program pacing” on page 223](#).

Model application program definitions

You can code a model application program definition, which can be used as the definition for one or more VTAM application programs. You code a model application program definition by placing wildcard characters in the name field of an APPL definition statement that defines characteristics for one or more application programs.

If VTAM cannot find an active definition for the application program associated with the ACB specified on an OPEN macroinstruction, it searches a list of model definitions to find a name that best matches the name of the application program associated with the ACB. If it finds a match, it uses that definition statement to define the application program. If it does not find a match, it returns an error condition indicator to the application program.

An application program that is built from a model application program definition is called a *dynamic application program*, a *clone application program*, or a *replicated application program*.

Model application program definitions enable you to reduce the number of application program definitions you must code in VTAMLST.

You can use model application program definitions in nonsysplex and sysplex environments. The value of using them, however, is most apparent in sysplex environments.

In a sysplex environment, model application program definitions can be contained in a single VTAMLST accessible to all VTAMs in the sysplex. An application program can open its ACB on any VTAM in the sysplex. In addition, an application program can be moved to any VTAM in the sysplex by closing its ACB on one VTAM and opening it on another.

Note: It is recommended that model application program definitions be used for application programs that intend to use the multinode persistent session (MNPS) function on an end node (EN), and it is required that model definitions be used for such applications on a network node (NN). Activating a nonmodel image of the MNPS application at an EN or NN, when the same application is active elsewhere (on another EN or NN), can lead to APPN search failures.

Overview

To code a model application program definition, code an APPL definition statement to define application program characteristics that you expect to be used by one or more VTAM application programs. Use wildcard characters in the name field of the APPL definition statement. You can use the following characters:

Asterisk (*)

Represents 0 or more unspecified characters.

Question mark (?)

Represents a single unspecified character.

When placing wildcard characters in the name fields of model application program definitions, you should have some idea of which dynamic application programs might be built from those model definitions. You will probably want to use a naming scheme that ensures that dynamic application programs are not accidentally built from model application program definitions; that is, that an application program does not open its ACB using a model definition that you did not intend for it to use. See [Table 39 on page 279](#)

for examples of model application program definition names and the names of the dynamic application programs that could be built from them.

Coding guidelines

Coding guidelines for model application program names depend on whether you define only the required name (in the label field of the APPL definition statement) or both the required name and the optional name (on the ACBNAME operand of the APPL definition statement).

When only the required name is defined

A model application program name, including wildcard characters, can be a maximum of eight characters in length. A question mark (?) can be used anywhere in the model application program name. An asterisk (*) can be used in the second to eighth characters of the model application program name. Model application program names must be unique across all resources known to VTAM (including model CDRSCs).

There is no defined limit on the number of dynamic application programs that can be created from one application program model definition, nor is there a defined limit on the number of model definitions that you can define.

Model definitions can be defined in any number of application program major nodes. Those model definitions can appear in an application major node along with conventionally defined APPL definition statements. The model definitions and conventionally defined APPL definition statements can appear in any order.

Table 39 on page 279 shows examples of model application program names.

Table 39. Sample model application program names	
Model application program name	Names that could match
APPL*	Any name that begins with "APPL," followed by zero to four additional valid characters in length. Examples: APPL APPL1 APPL1A APPL01A APPL1234
?*APPL	Any name that begins with at least one valid character, plus zero to three additional valid characters in length, followed by "APPL." Examples: A1APPL A1AAPPL A01AAPPL A123APPL
AP?L	Any four-character name that begins with "AP," followed by one valid character in length, and ends with "L." Examples: AP1L APPL AP\$L AP0L

Table 39. Sample model application program names (continued)

Model application program name	Names that could match
APPL?	Any five-character name that begins with "APPL," followed by one valid character in length. Examples: APPL1 APPLS APPL3 APPL0
A*APPL	Any name that begins with "A," followed by zero to three valid characters in length, and ends with "APPL." Examples: AAPPL A1APPL A01APPL A\$12APPL A123APPL

Note: You cannot specify a wildcard name that begins with an asterisk (*) because VTAM interprets this as a comment line. To achieve the effect of a wildcard name that begins with an asterisk, code two APPL definition statements:

- One whose name begins with ?* and ends with zero to six valid characters in length
- One whose name is the same as the zero to six characters

For example, *APPL would match any name that ended with APPL and began with zero to four valid characters in length. ?*APPL matches any name that begins with one to four valid characters in length, but it does not match a name that begins with zero to four characters in length, namely APPL. To allow for the possibility that an application program with a name ending in APPL and beginning with zero to four valid characters in length could open its ACB, you would have to code the following two APPL definition statements:

```
?*APPL  APPL  AUTH=(PASS,ACQ),EAS=5,...
APPL    APPL  AUTH=(PASS,ACQ),EAS=5,...
```

When the required name and the optional name (ACBNAME) are defined

When you use wildcard characters in the optional name specified on the ACBNAME operand, observe the following additional coding guidelines:

- At most, only one asterisk (*) can be used in the name specified on the ACBNAME operand. Up to eight question marks (?) can be used.
- The number of asterisks and question marks used in the name specified on the ACBNAME operand must match the number of asterisks and question marks used in the name specified in the label field of the same APPL definition statement.

Note that the asterisks and question marks do not need to be in the same positions in the required and optional names.

All other coding guidelines that apply when only the required name is defined also apply when both the required name and the optional name are defined.

When VTAM finds a match on an ACBNAME operand, the characters on the APPLID operand on the ACB macroinstruction matching the wildcard characters in the optional name specified on the ACBNAME operand are substituted in the required name, in left-to-right order, by wildcard type. That is, the character that matches the first question mark specified on the ACBNAME operand is substituted for

the first question mark in the required name. The character that matches the second question mark is substituted for the second question mark in the required name. All the characters (including none) in the name specified on the APPLID operand on the ACB macroinstruction that match the asterisk in the optional name specified on the ACBNAME operand are substituted for the asterisk in the required name.

How VTAM finds the best match

If VTAM finds more than one model application program definition that could match the name of an application program requesting to open its ACB, it uses the model definition whose name best matches the name of the application program requesting to open its ACB. (Similarly, if VTAM finds more than one model CDRSC definition that matches the name of a cross-domain resource it encounters, it uses the model definition whose name best matches the name of the cross-domain resource.) VTAM determines the best match following the rules for wildcard support described in [z/OS MVS Setting Up a Sysplex](#). This section provides an overview of these rules (which apply to model CDRSCs, and to model applications).

VTAM compares a model application program name with the name of the application program trying to open its ACB. It compares each name, character for character, scanning from left to right. For each character, VTAM looks for a match in the following order:

1. Most specific match: The actual character (A–Z, 0–9, @, #, or \$)
2. Next most specific match: A single-character wildcard character (?)
3. Least specific match: A multiple-character wildcard character (*)

This order implies the following situation:

- The model application program name that contains no wildcard characters will always be the most specific.

For example, if an application program named ABCDEFG requests to open its ACB, VTAM will search VTAMLST for a definition of ABCDEFG. If there is a definition statement named ABCDEFG, VTAM will use that statement as the definition of ABCDEFG.

- The model application program name with the most actual characters appearing before a wildcard character is more specific.

For example, if an application program named ABCDEFG requests to open its ACB, VTAM will search VTAMLST for a definition of ABCDEFG. If there is no definition statement named ABCDEFG, but there are two model definitions named ABC?EFG and ABCD*, VTAM will use ABCD* as the definition of ABCDEFG. As shown below, VTAM scans from left to right, looking first for character-to-character matches, which it finds for the first three characters in ABC?EFG. However, as it continues scanning, it finds that ABCD* has a character-to-character match for the first four characters, which gives it priority over ABC?EFG. Even though the question mark in ABC?EFG is in an earlier position (fourth) than the asterisk in ABCD* (fifth), because ABCD* matches ABCDEFG character-to-character for the first four positions, and because VTAM searches first for character-to-character matches, ABCD* is the more specific match.

```
A B C D E F G
| | |
A B C ? E F G
```

```
A B C D E F G
| | | |
A B C D *
```

Other examples are:

ABC?DE

is more specific than

ABC*DE

because the single-character wildcard character (?) is more specific than the multiple-character wildcard character (*).

ABC?DE

is more specific than

ABC??E

because the actual character (D) is more specific than the single-character wildcard character (?).

ABC?DE

is more specific than

ABC?D*

because the actual character (E) is more specific than the multiple-character wildcard character (*).

Tip: Use the DISPLAY MODELS command with the APPL operand to verify that the model definition that you intend to use for your application name is the one that VTAM selects.

Example of using model application program definitions

The following examples show how model application program definitions are used. One example shows how definitions are used when wildcard characters are used only in the required name. The other example shows how definitions are used when the ACBNAME operand is coded and wildcard characters are used in the optional name defined on the ACBNAME operand.

When only the required name is defined

The following model application program definitions are contained in VTAMLST:

```
APPL?  APPL  AUTH=(PASS,ACQ),EAS=5,...
APPL*  APPL  AUTH=(PASS,ACQ),EAS=1,...
A*APPL APPL  AUTH=(PASS,ACQ),EAS=6,...
AP?L   APPL  AUTH=(PASS,ACQ),EAS=3,...
```

The OPEN macroinstruction, OPENXYZ, specifies that the application program associated with ACB123 requests to open its ACB, as follows:

```
OPENXYZ  OPEN  ACB123
ACB123   ACB    AM=VTAM,APPLID=APPL1A,...
```

The ACB macroinstruction, ACB123, requests to be associated with the definition of application program, APPL1A. VTAM cannot find a definition for APPL1A in VTAMLST, but it does find a name that best matches APPL1A, and that is APPL*. The asterisk in APPL* indicates that it can be used as the definition for any application program whose name begins with APPL, followed by zero to four valid characters in length. APPL? is not a valid match in this example because the question mark indicates that it can be used as the definition statement for any application program whose name begins with APPL, followed by only one valid character in length.

When the required name and the optional name (ACBNAME) are defined

The following model application program definitions contain an optional name defined on the ACBNAME operand:

```
NET?NAME APPL  ACBNAME=WILD?,AUTH=(PASS,ACQ),EAS=5,...
NET*NM    APPL  ACBNAME=WILD*,AUTH=(PASS,ACQ),EAS=1,...
NET?NAM?  APPL  ACBNAME=WILD??,AUTH=(PASS,ACQ),EAS=6,...
```

The OPEN macroinstruction, OPENXYZ, specifies that the application program associated with ACB123 requests to open its ACB, as follows:


```
OPENXYZ  OPEN  ACB123
ACB123   ACB   AM=VTAM,APPLID=WILD1,...
```

The ACB macroinstruction, ACB123, requests to be associated with the definition of application program, WILD1. VTAM cannot find a definition for WILD1 in VTAMLST, but it does find a name, specified on the ACBNAME operand, that best matches WILD1, and that is WILD?. The ? in WILD? is replaced with 1, as is the ? in NET?NAME. Therefore, the optional name of the application program is WILD1 and the required name is NET1NAME.

If the name specified on the APPLID operand on the ACB macroinstruction was WILDONE, VTAM would use WILD* as the model application program definition. The optional name would be WILDONE and the required name would be NETONENM.

If the name specified on the APPLID operand on the ACB macroinstruction was WILD1S, VTAM would use WILD?? as the model application program definition. The optional name would be WILD1S and the required name would be NET1NAMS.

Resource state requirements

A model application program must be active before it can be used to build dynamic application programs. You can activate a model application program as you activate other application programs by doing one of the following actions:

- Issue a VARY ACT command.
- Include the major node in which the model application program is defined in the configuration list that VTAM uses when it is initialized.

By activating the model application program, you ensure that the state of the model application program is connectable, thus making it available to build dynamic application programs.

You cannot activate dynamic application programs with the VARY ACT command. Dynamic application programs are activated only by opening their ACBs. For a dynamic application program to successfully open its ACB, a model application must exist to build the dynamic application program, and that model application program must be connectable.

Authorizing application facilities

You can use the following facilities to enhance your application program environment:

- Passing and validating logon requests
- Overriding dial number digits
- Acquiring LU sessions
- Enabling parallel sessions
- Authorizing privileged paths

Passing and validating logon requests

You can write a logon exit routine that passes VTAM session requests to another application program. This allows the application program to validate the logon before accepting it. Passing an LU session request to another application program queues the LU session request to that application program. A session is established after the original session has ended. The AUTH=PASS operand on the APPL definition statement permits the passing of session requests.

Overriding dial number digits for dial connections

In a dial-out operation, an application program can be authorized to override the dial number digits and other parameters specified in the switched major node for the device being contacted. The device can be PU types 1, 2, or 2.1. Normally, if these parameters do not match in XID (exchange identification), the connection is terminated. However, you can authorize an application to override this checking. Coding the

AUTH=ASDP operand on the APPL definition statement allows the application to supply dial parameters to the model switched major node during session initiation.

Notes:

1. This operation can cause a security exposure.
2. If an active APPN connection exists, VTAM may attempt to use it rather than dialing a new connection.

Acquiring LU sessions

An application program acquires LU sessions by issuing either an OPNDST macroinstruction with OPTCD=ACQUIRE or a SIMLOGON macroinstruction followed by an OPNDST macroinstruction with OPTCD=ACCEPT specified.

You can enable the acquiring of sessions by coding the AUTH=ACQ operand on the APPL definition statement.

Enabling parallel sessions

Certain session protocols can prevent an application program from performing more than one transaction at a time on a single session. Using the parallel sessions function, you can enable an application program to engage in several sessions at once with the same logical unit. Thus, the application program can perform multiple transactions at the same time.

If you want an application program to have parallel sessions, you must authorize it by coding PARSESS=YES on the APPL definition statement. APPC=YES overrides PARSESS=NO to allow parallel sessions. The application program can distinguish sessions between the same LUs through the proper use of the communication identifier (CID).

Authorizing privileged paths

Authorized path is a VTAM facility that enables an application program to specify that a data transfer operates in a privileged manner. Authorized path enables an authorized application program to run entirely under a system request block (SRB) after the application program ACB has been opened. In most cases, this means a shorter path length through VTAM for RPL-related macroinstructions. An application program can be authorized by using the authorized program facility (APF) running in the supervisor state or by running in a key of 0 through 7. For a full description of authorized path, see [z/OS Communications Server: SNA Programming](#).

Data compression

When using application programs in a multiple-domain environment, consider taking advantage of the data compression facility. This facility enables VTAM to compress the data for messages on selected LU-LU sessions. Data compression is transparent to application programs.

When using application programs in a single-domain environment, between application-to-application sessions in the same host, there is no benefit to compressing data. In fact, extra processing cycles are expended. If the SLU is an application in the same host, VTAM builds the BIND without data compression even if data compression is indicated. Data compression is not allowed in these single-domain conditions.

Also, you should not use the data compression facility for VTAM in conjunction with any application program data compression. After the application program has compressed the data, the amount of additional compression would be minimal.

You can implement data compression for logical units in either an extended BIND (types 2.1, 4, or 5 physical unit) or nonextended BIND environment (type 1 or type 2 physical unit). You cannot compress data for sessions with LU type 4 or LU type 7 in a nonextended BIND environment.

Data compression uses data storage.

Types of compression

The use of data compression can be negotiated for each session through the BIND and BIND(RSP) request units (RUs). The session partners can negotiate whether to use data compression. If compression is to be used, the session partners can negotiate the type of compression. VTAM supports two different compression algorithms:

- Run length encoding (RLE) compression

This type of compression replaces strings of identical bytes with shorter encoded strings.

- Dynamic dictionary-based compression (called adaptive compression in this document)

This type of compression is an adaptive dictionary-based compression algorithm similar to Lempel-Ziv. It is applicable to a wide range of data types. It uses tables that adapt dynamically to match the data being sent or received. Adaptive compression replaces the original data with a set of compression codes. Each set represents one or more bytes. Following are the variations of adaptive compression:

Small table

This type uses 9-bit compression codes.

Medium table

This type uses 10-bit compression codes.

Large table

This type uses 12-bit compression codes.

With adaptive compression, each partner begins a session with an identical set of tables. The senders of the RU update their send tables as data is compressed. The receivers of the RU makes identical updates to their receive tables as data is extracted. This method keeps both ends of the session identical without exchanging table data between nodes.

By default, the send and receive tables are continually updated. You can speed up adaptive compression processing on outbound messages, through software only or with the help of hardware compression, by using the CMPMIPS start option. By specifying a value from 1 to 99 (default value is 100), VTAM suspends and resumes updating of the send and receive tables whenever compression effectiveness exceeds or falls below a threshold corresponding to the CMPMIPS start option value.

Thus, using the CMPMIPS option, you can balance the number of machine cycles needed with the effectiveness of compression for outbound messages. Higher values for CMPMIPS will likely increase both compression effectiveness and cycle usage, while lower CMPMIPS values will likely lower both compression effectiveness and cycle usage.

Implementing data compression

If you choose to use data compression, you can adjust the use of compression for your system by using a combination of compression levels specified for VTAM and for application programs. The following levels of data compression are available:

0

No compression

1

RLE

2

Adaptive small table

3

Adaptive medium table

4

Adaptive large table

- The overall compression limit for a VTAM host can be:

- Initialized through the CMPVTAM start option. For more information about the CMPVTAM start option, see [z/OS Communications Server: SNA Resource Definition Reference](#).

- Changed by the MODIFY COMPRESS command. (The compression levels of active sessions are not affected.) For more information about this command, see [z/OS Communications Server: SNA Operation](#).
- Displayed through the DISPLAY VTAMOPTS command. For more information about this command, see [z/OS Communications Server: SNA Operation](#).
- Input and output compression limits for an application program can be:
 - Initialized through the CMPAPPLI and CMPAPPLO operands on the APPL definition statement.
 - Changed by the MODIFY COMPRESS command. (The compression levels of active sessions are not affected.)
 - Displayed through the DISPLAY ID command.
- COMPRES is the compression-override operand on the MODEENT macroinstruction that specifies how the negotiation of compression should be handled. You can use COMPRES to override the BIND negotiation process and either require or prohibit compression. Using this operand data compression specifications, you can associate with a logon mode name. For more information about this operand, see [z/OS Communications Server: SNA Resource Definition Reference](#).
- The DISPLAY SESSIONS command can indicate the compression levels in use for a given session and the average message length reductions achieved in both directions. For more information about this command, see [z/OS Communications Server: SNA Operation](#).

Compression level negotiation

If the value specified on the COMPRES operand of the MODEENT macroinstruction allows negotiation, the compression levels for a session are negotiated through the BIND and BIND(RSP) RUs.

The data compression information is carried on a control vector X'66'. Control vector X'66' is used only for compression-level negotiation. The control vector is composed of these subvectors:

X'80' is passed on a CDCINIT and a CINIT.

X'81' is passed on a BIND to carry the requested compression levels.

X'82' is passed on a BIND(RSP) to carry the actual compression levels.

The SLU first requests a compression level on the CDCINIT. The PLU responds with requested levels specified on the BIND. The SLU responds with actual levels on the BIND(RSP).

Note: There are limited negotiation capabilities for sessions using a nonextended BIND. Bits 6 and 7 in byte 25 of the BIND are reserved for data compression, and are used to negotiate whether compression is acceptable or not. No level negotiation is possible. In the PLU-to-SLU direction, level 2 is the highest possible compression level (If VTAM owns the PLU, setting CMPVTAM = 1 further limits the PLU-to-SLU level to 1). In the SLU-to-PLU direction, level 1 is the highest possible compression level.

Negotiation when VTAM owns the PLU

When VTAM owns the PLU for the session, VTAM has the option to request compression. VTAM does not request compression if any of the following conditions occur:

- The CMPVTAM start option is set to 0, which indicates no compression.
- The CMPAPPLI and CMPAPPLO operands on the APPL definition statement are set to 0, which indicates no compression.
- The SLU owner prohibits compression through bits in the control vector X'66'. Whether the control vector indicates that compression is prohibited or required is determined by the setting of the COMPRES operand of the MODEENT macroinstruction.

If none of the previous conditions occur, the PLU requests compression levels. The PLU-to-SLU compression level is the lower of the values set on the CMPVTAM start option or the CMPAPPLO operand on the APPL definition statement. The SLU-to-PLU compression level is the lower of the values set on the CMPVTAM start option or the CMPAPPLI operand on the APPL definition statement.

Negotiation when VTAM owns the SLU

When VTAM owns the SLU for the session, the SLU can begin negotiation when it sends the control vector X'66'. The control vector indicates whether compression is prohibited or required. If compression is required, the COMPRES operand on the MODEENT macroinstruction must indicate that compression is required.

When sending the BIND(RSP), the SLU negotiates no compression if any of the following conditions occur:

- The BIND does not request compression.
- PROHIB or SYSTEM was specified in the control vector X'66'.
- The CMPVTAM start option is set to 0, which indicates no compression.

Note: When compression is negotiated from the SLU side, the CMPAPPLI and CMPAPPLO operands are not used.

If none of the previous conditions occur, the SLU requests compression levels. The PLU-to-SLU compression level is the lower of the values set on the CMPVTAM start option or the requested PLU-to-SLU compression level passed on the BIND. The SLU-to-PLU compression level is the lower of the values set on the CMPVTAM start option or the requested SLU-to-PLU compression level passed on the BIND.

Examples of data compression

The following are examples of how the values set for CMPVTAM, CMPAPPLI, and CMPAPPLO work together to determine the compression levels for a session. [Figure 87 on page 287](#) shows a sample configuration.




Figure 87. Data compression yield

For the examples, assume the values shown in [Table 40 on page 287](#). Also assume that COMPRES=REQD is specified on the logon modes used for the sessions.

Table 40. Compression values for example of data compression yield			
VTAM1	APPL1	APPL2	VTAM2
CMPVTAM = 4	CMPAPPLI = 4	CMPAPPLI = 4	CMPVTAM = 3
	CMPAPPLO = 1	CMPAPPLO = 2	

In the first example, APPL1 is the PLU and APPL2 is the SLU.

1. VTAM determines that the requested PLU-to-SLU level (which is the lower of the values specified on the CMPVTAM start option for VTAM1 or the CMPAPPLO operand for APPL1) is 1.
2. VTAM determines that the requested SLU-to-PLU level (which is the lower of the values specified on the CMPVTAM start option for VTAM1 or the CMPAPPLI operand for APPL1) is 4.
3. APPL1 passes the requested levels to APPL2 on the BIND.
4. VTAM compares the requested levels to the VTAM2 CMPVTAM level. APPL2 determines that the actual PLU-to-SLU level for this session is 1 and the actual SLU-to-PLU level for this session is 3.
5. APPL2 returns the actual levels on the BIND(RSP).

In the second example, APPL2 is the PLU and APPL1 is the SLU.

1. VTAM determines that the requested PLU-to-SLU level (which is the lower of the values specified on the CMPVTAM start option for VTAM2 or the CMPAPPLO operand for APPL2) is 2.
2. VTAM determines that the requested SLU-to-PLU level (which is the lower of the values specified on the CMPVTAM start option for VTAM2 or the CMPAPPLI operand for APPL2) is 3.
3. APPL2 passes the requested levels to APPL1 on the BIND.
4. VTAM compares the requested levels to the VTAM1 CMPVTAM level. APPL1 determines that the actual PLU-to-SLU level for this session is 2 and the actual SLU-to-PLU level for this session is 3.
5. APPL1 returns the actual levels on the BIND(RSP).

Summary of data compression

The following summarizes data compression.

- Data compression is established between users of a session through the BIND and RSP(BIND) flows. VTAM can handle compression settings at one end or both ends on behalf of VTAM applications.
- Compression is disabled on sessions with both ends in the same VTAM.
- VTAM defines compression levels 0 to 4 used in VTAM start option CMPVTAM and PLU-APPLication definition options CMPAPPLI (incoming flow) and CMPAPPLO (outgoing flow). The actual compression level settles on the lower of the CMPAPPLI value or the CMPAPPLO value, and the level acceptable for each of the hosts (CMPAPPL value for VTAM).
- The CMPMIPS option can be used to balance CPU cycles versus LZ-like compression efficiency for outgoing data only. The value (0–100) can be changed dynamically using the command

```
F NET,VTAMOPTS,CMPMIPS=value
```

When CMPMIPS=0, data on all sessions is sent uncompressed. When CMPMIPS=100, most CPU cycles are used for best compression efficiency. A recommended value of 50 should save CPU cycles while maintaining high compression efficiency.

- RLE compression is allowed with its normal CPU utilization anytime CMPMIPS is not 0.
- Compression efficiency may be displayed on selected sessions using the command

```
D NET,SESSIONS,SID=sessionid
```

Efficiency is shown as a percentage reduction value on the original length of the uncompressed data.

- Some VTAM compression level statistics may be displayed using the command

```
D NET,STATS,TYPE=COMPRESS
```

- The logmode table option COMPRES= SYSTEM | PROHIB | REQD can be used for the logmode table entry for the sessions over which compression is requested.
 - The logmode table applies to the SLU-controlling VTAM.
 - The default value SYSTEM may be OK when the remote station is smart enough to accept, reject or update the compression levels requested by VTAM on the BIND CV66. This may be the case for a CM/2 acting as node 2.1.
 - SYSTEM=REQD is required for VTAM APPL to VTAM APPL sessions. For a work station with compression capability but no compression negotiation capability on CV66, set COMPRES=REQD. The PLU VTAM compression levels (resolved from CCMPVTAM, CMPAPPLI, or CMPAPPLO) must be acceptable to the work station.

Security features

VTAM provides a variety of security features that can be enabled for an application program, including:

- Cryptography facility

- Message authentication
- SLU authentication
- VTAM application security
- Confidential data
- 3270 Intrusion Detection Services

For information about LU 6.2 security features, see [“LU 6.2 security”](#) on page 334.

Cryptography facility

The cryptography facility protects the confidentiality of data transmitted between network resources by enciphering and deciphering session data. Cryptography is available for both LU 6.2 and non-LU 6.2 sessions. Support is available for both switched and nonswitched LUs. However, support is not available for binary synchronous communication (BSC) or local non-SNA devices.

The facility establishes cryptographic sessions for application programs and peripheral node LUs that require cryptographic services. For an LU to have a cryptographic session, the host processor must support cryptography.

The encryption facility provides two levels of cryptographic sessions:

Selective

Each end of the session selects the data to be enciphered before transmission. The selection is based on the capability of the session partner and the availability of cryptographic services.

Required

All outbound data request units are enciphered and all inbound data request units are deciphered.

The encryption facility uses services provided by the z/OS Integrated Cryptographic Service Facility (ICSF) and IBM Z Cryptographic Co-Processor. ICSF is a licensed program that runs under MVS and provides access to the hardware cryptographic feature for programming applications. The combination of the hardware cryptographic feature and ICSF provides secure high-speed cryptographic services.

Tip: You might be able to use cryptographic products other than ICSF if the cryptographic product runs in one of the supported modes of operation. The following terms are used to see cryptographic products that support one of these modes of operation:

- PCF/CUSP - Refers to any cryptographic product that is compatible with PCF/CUSP.
- CCA - Refers to any cryptographic product that is compatible with Common Cryptographic Architecture (CCA).

Requirement: Triple-DES 24-byte encryption requires the use of Common Cryptographic Architecture (CCA). CCA defines a set of cryptographic functions, external interfaces, and a set of key management rules that provide a consistent, end-to-end cryptographic architecture across different IBM platforms.

The cryptographic facilities provide services that include handling requests that VTAM receives to generate a cryptographic key. The cryptographic key is used to encipher and decipher session data.

Using dynamic cryptographic keys, you can do the following actions:

- Define both unique and alternate key-encrypting key names for LUs and CP/SSCPs.
- Switch between cryptographic products while VTAM is running.

Note: Switching to PCF/CUSP will terminate any sessions using triple-DES.

- Establish "clear" sessions (without encryption) if ENCR=COND and when either session partner does not support cryptography, or when cryptographic services are temporarily unavailable.

Implementing cryptography for LU-LU session data

Procedure

Take the following steps to use the cryptography feature:

1. Install and activate the cryptographic product.

The ENCRYPTN start option enables you to start VTAM before activating the cryptographic product. (ENCRYPTN=CCA is required if triple-DES encryption will be used by LUs or applications owned by this node.)

2. File cryptographic keys.

File the cryptographic key data set at the host processor before activating any LUs that are used in cryptographic sessions.

For information about filing cryptographic keys, see [Appendix E, “Cryptographic keys,” on page 593](#).

3. Specify the cryptographic requirements.

VTAM uses this information to identify the cryptographic session keys and to establish the cryptographic session.

- Code the ENCR operand and the ENCRTYPE keyword (DES or TDES24) on the LU and APPL definition statements to define the cryptographic capabilities of LUs and application programs.
- Code the ENCR operand on the MODEENT macroinstruction to specify cryptographic session requirements for an LU.
- Code the ENCR TYP operand on the MODEENT macroinstruction to specify encryption type for an LU (TDES24 is the only available value).
- Identify the name of the cryptographic key that will be used to establish cryptographic sessions for the LU. Code the CKEYNAME operand on the LU definition statement or use the default, which is the LU name.
- Specify whether you want to use the alternate cryptographic key name during session activation. You can either use the CKEY operand on the MODEENT macroinstruction or issue a modify security command to switch to the alternate CKEY.

4. Initiate the cryptographic session.

Note: For information about the options you can specify on the ENCR, ENCRTYPE, CKEYNAME, and CKEY operands, see [z/OS Communications Server: SNA Resource Definition Reference](#).

Establishing cryptographic sessions

Cryptographic LU-LU sessions are established based on the requirements and capabilities of the session partners.

As shown in [Figure 88 on page 291](#), the application programs can establish the following sessions with the peripheral node LUs:

APPL1A

Can establish a nonencrypted session with LU1A.

Can establish a cryptographic session with LU1B *if* a cryptographic product is currently available. If it is not available, a session without encryption is started.

APPL1B

Can establish a cryptographic session with LU1B *if* a cryptographic product is currently available; otherwise the session fails.

Cannot establish a session with LU1A; the session fails. APPL1B requires a cryptographic session that LU1A does not support.




Figure 88. Encryption facility specifications

The MODIFY ENCR or MODIFY SECURITY can be used to change the cryptographic capability of a logical unit, or change the capability to a higher level of cryptography than was previously specified. Only MODIFY SECURITY can be used to change the minimum encryption type (ENCRTYPE).

When VTAM tracing is in effect for a cryptographic session, the protected data in the VTAM buffers is neither displayed nor printed.

When sessions require encryption, a session could fail if a cryptographic support is not currently available. You can define some sessions to establish a clear session if cryptographic services are not available at session start. The ENCR=COND operand on the APPL definition statement specifies whether to establish a clear session if a cryptographic facility is not available. A clear session is established when a cryptographic facility is not available and one of the following conditions exist:

- Both session partners have ENCR=COND coded.
- When ENCR=COND is coded by one session partner and ENCR=OPT is coded for the other session partner.

If either end of a session specifies a higher level of encryption than ENCR=COND (ENCR=SEL or ENCR=REQ), an encrypted session is required and the session request fails. The following table shows examples of the interactions of the ENCR and ENCRTYPE keywords in determining the encryption level.

Table 41. DES-TDES24 encryption options				
ENCR		ENCRTYPE		Results
SLU	PLU	SLU	PLU	
REQD	REQD	DES	DES	8-byte DES encryption is used.
REQD	REQD	TDES24	DES	TDES24 encryption is used if PLU hardware is capable of supporting triple-DES. The PLU side can be upgraded to triple-DES. Encryption fails if PLU hardware is not capable of supporting triple-DES.
REQD	REQD	DES	TDES24	Encryption fails; the SLU has obtained too small a key (8 bytes) while the PLU requires a 24-byte key. The SLU VTAM cannot be upgraded; it obtains the session key.
OPT/REQD	REQD	DES*	TDES24	(* — SLU also sets &ENCRYP=TDES24) TDES24 encryption used, assuming both SLU and PLU hardware is capable of supporting triple-DES.
OPT	COND	DES	TDES24	Session established in the clear. The ENCRTYPE mismatch cannot be resolved, but encryption is not required.

Information needed to encipher and decipher session information is included in session establishment commands. When a session that will use cryptography is being initiated, the cryptographic facility of the SLU SSCP/CP enciphers the cryptography key. From the cryptography key two enciphered keys are created:

1. One copy is enciphered under the SLU master key

The enciphered session key is placed in the BIND image to be used later. For 24-byte keys, the extra 16 bytes are transported as control vector data and not in the BIND image itself.

2. Another copy is enciphered under a cross key.

The enciphered cross key is inserted in the CDINIT or CDCINIT command (or corresponding APPN command) and is used by nodes along the session path. Depending on the configuration, this key is deciphered and reenciphered at every node along the session path or only by the session endpoint.

See [Appendix E, “Cryptographic keys,”](#) on page 593 for a description of a cross key and how these keys are entered into the various cryptographic products.

Note: The actual session data that is enciphered under the session key is enciphered and deciphered only at the session endpoints.

Cryptographic session initiation

The following only applies to the session key *enciphered under a cross key* as it is used during session initiation.

VTAM supports both end-to-end and host-by-host encryption. The method used is dependent on the types of nodes in the configuration and the coding of the cryptographic key data sets (CKDS).

End-to-end

Session key is deciphered and or reenciphered only at the session endpoints or APPN endpoints during session initiation.

Host-by-host

Session key is deciphered and or reenciphered at every VTAM along the path during session initiation.

The installation actually determines which method VTAM will use by placing cross keys in the appropriate cryptographic facility data sets and by the capability of the nodes involved.

- If VTAM is defined as an APPN node and, if during session setup VTAM finds the name of the CP or network node server (NNS) of the target, that VTAM can provide end-to-end cryptography support for subarea, APPN, and mixed subarea and APPN networks.
- If VTAM is defined as a subarea only node, the session cryptography key usually must be encrypted on a host-by-host basis.

However, if the VTAM node that is defined as a subarea only is connected to an NN, the first host-by-host decipher and reencipher will be done from the subarea only VTAM to the NN. When the NN (by definition an APPN-capable node) chooses the cross key for the next leg of the session initiation, it may find the name of the CP of the target so this host-by-host decipher and reencipher will be altered to an end-to-end decipher and reencipher.

During session initiation VTAM interrogates the encryption facility to determine whether cross key has been defined for a particular name.

The following information describes the order in which VTAM will choose a name by which to interrogate the cryptographic product:

1. If the name is present, the cross key associated with that name will be used to encipher the session key for the next, or only, hop.
2. If the name is not present, the next favored choice of name will be selected and the interrogation will again be attempted.
3. If the entire list is attempted and no name is found, the attempt to initiate a cryptographic session will fail.

When preparing to send a session initiation request into an APPN network or when sending a subarea CDINIT (request or response):

1. VTAM chooses the name of the owning CP of the PLU.
2. If step one fails, and the information is known to VTAM, VTAM chooses the name of the NNS of the PLU. This choice is not possible for a subarea CDINIT because there is no network node server (NNS) involved in a subarea-only configuration.
3. If that fails, VTAM chooses the name of the adjacent node, but only if the adjacent node is either another VTAM using an SSCP-SSCP control session or an NNS for the session. Note that cryptographic processing cannot be done by intermediate APPN nodes because they do not parse the session initiation request.

When sending cryptographic information about a CDCINIT into a subarea network, VTAM always follows the earlier host-by-host algorithm — encrypt the key in the cross domain key of the adjacent node. Also, if cryptographic processing was done on CDINIT, it will not be done again on CDCINIT.

The first key found using the above search will be the key used. VTAM also includes the partner name in a control vector so the other VTAMs along the path either ignore the cryptographic fields when the name included is not theirs, or decipher the cross key. If VTAM deciphers the cross key, VTAM then acts upon the cross key by saving the key if this is the endpoint, or reenciphering the key, changing the name in the control vector, and then forwarding the session initiation.

End-to-end cryptography

End-to-end cryptography enables an LU-LU session to use encryption as long as all intermediate VTAMs are Version 3 Release 2 or later and both endpoint VTAMs are Version 4 Release 1 or later. Instead of filing cryptographic key data sets at each host that an LU-LU session initiation request traverses, the installation

defines the data sets at only the endpoint hosts. These data sets contain the cross keys for the other endpoint and not the adjacent node. Defining cryptographic key data sets between only endpoint hosts eliminates the need to share cryptographic key information with other hosts. This can be important if, for example, the intermediate hosts are owned by a different company. Using end-to-end cryptography in [Figure 90 on page 295](#), cryptographic key data sets would need to be filed in only VTAM1 and VTAM3.

Defining cryptographic keys between only endpoint hosts can still become burdensome, because many endpoints can exist and the amount of key definition can get very large. Therefore, in an APPN network, cryptographic key data sets can be defined in two ways:

- Between a network node server and its end nodes
- Between network node servers

Although more key definition is required at a network node server using this strategy, the overall number of cryptographic key data sets declines. End nodes are required to define only one key to the network node server (rather than a key for each node the end node is to communicate with). In addition to the keys for its served end nodes, network node servers are required to define only keys for other network node servers (rather than a key for each node its end nodes are to communicate with).

Both session endpoints do not have to use their network node servers for cryptography. That is, one end node can define a key to its server and the server can define a key for the destination end node (instead of defining a key for the destination end node network node server).

An example of defining cryptographic keys using the network node server is shown in [Figure 89 on page 294](#). If VTAM1 files a cryptographic key data set with VTAM2 and VTAM2 files cryptographic key data sets with VTAM4, applications and LUs at VTAM1 can have cryptographic sessions with applications or LUs at VTAM4 without the need for VTAM1 and VTAM4 to exchange keys.




Figure 89. Encryption facility in an APPN environment

Host-by-host cryptography

With encryption on a host-by-host basis, each host must support cryptography, and cryptographic key data sets must be filed at each host involved in the LU-LU session initiation.

In [Figure 90 on page 295](#), for cryptographic sessions to exist between resources in VTAM1 and VTAM3, cryptographic key data sets must be filed at VTAM1, VTAM2, and VTAM3.

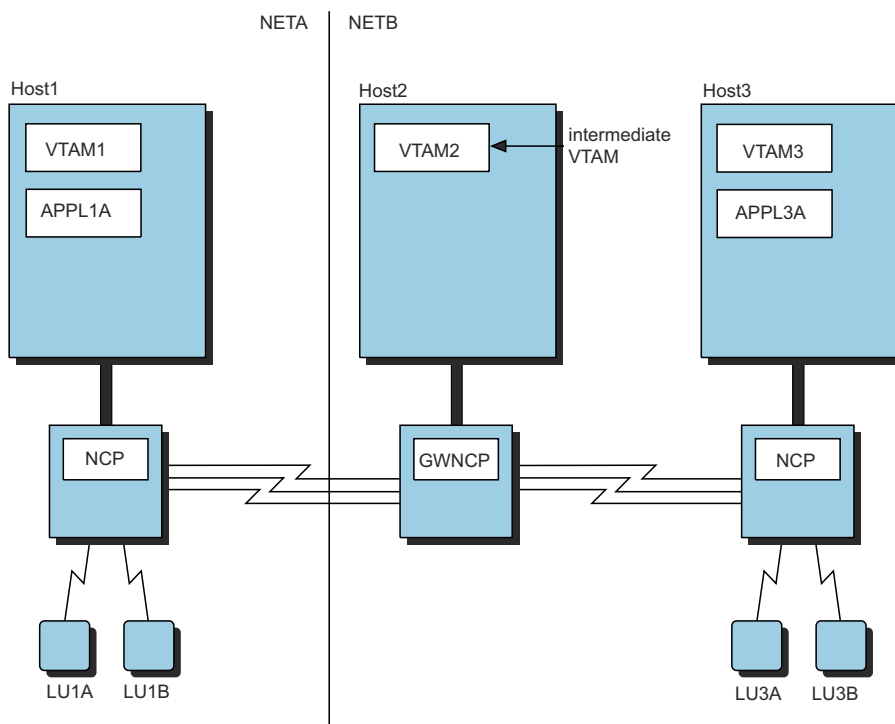


Figure 90. Encryption facility in multiple-network environment

Changing cryptographic keys dynamically

You can change LU and CP/SSCP resource keys for a session while VTAM is running. This enables you to avoid synchronization problems when changing keys, and helps reduce down time when changing static definitions.

By coding the appropriate operands and issuing the appropriate commands, you can:

- Use alternate LU key-encrypting keys while you change the LU master key in the CKDS at the control point (CP/SSCP).
- Use alternate CP/SSCP import and export key-encrypting keys (KEKs) while you change the corresponding import/export values in the CKDS at each control point.

Using alternate LU key-encrypting key names

To encipher the session key that is sent to the LU when a session is established, the control point (CP or SSCP) that owns the LU uses the LU master key in the cryptographic key data set (CKDS). When changing the LU master key, a similar change must also be made in the CKDS at the control point. Otherwise, the LU would have the new key and the CP would have the old key.

By specifying the CKEY operand VTAM uses an alternate key-encrypting-key (KEK) name that matches the LU master key while the primary master key is being changed in the CKDS. The CKEY operand can be specified on the MODEENT macro in the logon mode table or on the MODIFY SECURITY command. The KEK specified by the CKEY operand will be used during session activation until it is changed by the MODIFY SECURITY command or the LU is deactivated.

Use of the CKEY operand eases the administration process involved when changing the LU master key. Without the CKEY, both the master key at the LU and in the CKDS must be done simultaneously. This in

turn reduces the likeliness of a session failure because the LU master key and CKDS key are not the same. The CKEY allows for continued use of the alternate master key until both the CKDS and LU have been updated with the new primary master key.

For VTAM to use the alternate LU KEK name while you change the LU master key, follow this general procedure:

1. Notify user of a specific time when the value of the default KEK for the LU will be changed.
2. At the specified time, change the default and alternate KEK in the CKDS to the appropriate values.

Note: These changes only apply to subsequent sessions with the LU. Currently active or pending sessions use the old key values.

See [z/OS Communications Server: SNA Resource Definition Reference](#) for details on the CKEY operand, and to [z/OS Communications Server: SNA Operation](#) for full details on the use of the MODIFY SECURITY command.

Specifying key encrypting key (KEK) names for LUs

You can specify a cryptographic key name to specify a unique cryptographic key name for an LU by specifying the CKEYNAME operand. If you do not specify CKEYNAME, the cryptographic key name defaults to the LU name. You can code the same value for multiple LUs on the CKEYNAME operand, which reduces the number of definitions needed in the cryptographic key data set.

The CKEYNAME operand can be coded on the LU definition statement in the following major nodes:

- Local SNA
- LU group
- Model
- NCP
- Switched

Using import and export CP/SSCP KEK names

By using unique import and export KEKs, VTAM may support other cryptographic products that implement CCA and provide you with a choice between ICSF and these other cryptographic products. VTAM appends a unique prefix and suffix to the CP/SSCP name that is used to reference import and export KEKs.

To change the export and import KEKs in the CKDS of a host, you must also change these keys in the other host. Instead of bringing down sessions to change the master keys, you can force VTAM to temporarily use an alternate name that matches the new LU master key name. The user should update the LU master key in the CKDS at the CP/SSCP as soon as possible after being notified that the keys are changed.

Note: For migration purposes, VTAM tries the CP/SSCP name again without the suffix if a request fails because the KEK could not be found for the CP/SSCP name with the suffix. This alleviates having to change the CKDS.

Follow these steps when using alternate KEK names for CPs and SSCPs:

1. File an alternate set of export and import keys in each CKDS, in addition to the original export and import key-encrypting key names. You must file keys in the CKDS of each host on each end of a cross-domain session or each end of an APPN session (for example, HOST1, HOST2):

HOST1

- IMPORTER.CP2
- EXPORTER.CP2
- IMPORTER.CP2.ALT
- EXPORTER.CP2.ALT

HOST2

- IMPORTER.CP1

- EXPORTER.CP1
 - IMPORTER.CP1.ALT
 - EXPORTER.CP1.ALT
2. Delete export keys in each host.
 - HOST1**
EXPORTER.CP2
 - HOST2**
EXPORTER.CP1
 3. Issue the MODIFY SECURITY command to specify that VTAM use the .ALT KEKs.
 4. Define new import and export values for the KEK fields in the CKDS of each host. VTAM starts all new sessions using the new values. Keep the alternate names in the CKDS until you are sure that all pending sessions using the alternate names have completed.

Message authentication

Message authentication is another form of security. Similar to data encryption to ensure data confidentiality, the message authentication data security feature:

- Provides services to ensure the integrity of data for selected LU-LU sessions.
- Provides end-to-end protection of data, which does not require support from intermediate nodes.

Message authentication allows VTAM to determine if a message has been altered in transmission between the session partners. A code is attached to each message by the sender and verified by the session partner.

There are two methods for producing the message authentication code:

- Data encryption standard (DES) product that requires a cryptographic product to be active. Using this method, both cryptography and message authentication can be performed concurrently. Although the keyword is DES, if the session is setup to use triple-DES encryption, TDES24 will be used. The use of the term DES here does not mean only DES encryption can be used.
- Cyclic redundancy check (CRC), which creates a message authentication code using an internal VTAM algorithm. Using this method does not require a cryptography product to be active.

The APPL definition statement and MODEENT macroinstruction provide operands that you can use to define the message authentication support to be provided for a session. Code the following operands for each end of the session:

MAC

Specifies whether authentication of data sent and received by the LU is required, conditional, or not supported.

MACLNTH

Specifies the minimum length of the message authentication code attached to the message.

MACTYPE

Specifies the type of message authentication checking (DES or CRC) to be used for the session.

See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about how to code the preceding operands.

Modifying and displaying MAC specifications

Using the MODIFY SECURITY and the DISPLAY ID commands, you can display or change the message authentication specifications for an LU.

MODIFY SECURITY

Increase the security specifications as defined in the applications' resource definition. Any change you make must increase the level of the message authentication required for the application.

DISPLAY ID

Display the current security data for the application. The DISPLAY ID has session status modifiers to show what a given session level is using for the level of encryption (DES = D or Triple-DES = T) encryption. See [z/OS Communications Server: IP and SNA Codes](#).

See [z/OS Communications Server: SNA Operation](#) for details on using the preceding commands.

SLU authentication

During establishment of a cryptographic session, the secondary LU (SLU) verifies that the primary LU (PLU) is using the same cryptographic key. However, if the SLU does not reject a session when the partners are using different session keys, subsequent data flowing over the session is lost because proper deciphering of the data cannot be performed by the SLU. By specifying the CERTIFY=YES operand, the verification that both partners are using the same session key is performed by both the PLU and the SLU. If the session keys are not the same, the session is not established.

The CERTIFY operand can be coded for a resource in the following major nodes:

- Application program
- Local SNA
- Logical unit (LU) group
- Model
- Network Control Program
- Switched

In addition, the CERTIFY operand can be specified on the MODEENT macro in the logon mode table. See the [z/OS Communications Server: SNA Resource Definition Reference](#) for additional information about coding the CERTIFY operand.

Notes:

1. SLU authentication is performed if CERTIFY=YES is specified for either session partner.
2. CERTIFY=YES is supported by both the 3174 and 3276 cryptographic features. To determine support for other devices, see the specific device documentation.

VTAM application security

You can specify a password on the APPL definition statement and require the application program to specify both its APPL definition statement name and its password when it opens its access method control block (ACB). The authority of the application program to gain access to the network can be verified by comparing the password in the OPEN ACB macroinstruction of the program to the password specified on the PRTCT operand of the APPL definition statement.

The VTAM application OPEN ACB security facility provides additional resource access control. To perform security processing, VTAM invokes a security management product (such as RACF) through the security authorization facility for any application program that is not APF-authorized. The security management product determines whether an application program access to the network is approved.

To enable the VTAM application OPEN ACB security facility, register the application program with a class of VTAMAPPL (CLASS=VTAMAPPL) in a security management product that is capable of controlling authorization for VTAM application program execution (such as RACF Version 1 Release 9). VTAM bypasses any password checking if the security management product provides the resource access control. The application program is authorized to access the network based on either the approval of the security management product or the user-specified password check.

The authority of the LU to initiate a session with an application program can also be verified by requiring that the LU include a password in its logon. This password can be verified by the session authorization function of a session management exit routine, a session authorization exit routine, the application logon exit routine of the program, or the application program itself.

Confidential data

When data is transmitted between an application program and a logical unit, it passes through VTAM buffers. These buffers are allocated from common page-fixed buffer pools when I/O is being performed and from pageable buffers in user-protected storage when data is queued. An application program that is transmitting or receiving confidential data can have fixed buffers cleared if PROC=CONFTXT is specified in the appropriate node initialization block (NIB). If you use this option, only the name of the application program, the name of the logical unit, and the direction of the data flow are included in the trace records; confidential data is not included. A buffer trace of nonconfidential data includes the data.

Following are two ways of protecting application program data:

- The application program can ensure that confidential data in VTAM buffers within the host of the application program is cleared after the data is sent to the NCP or is moved into the application program address space. The application program does this by using the CONFTXT option in the node initialization block (NIB) associated with the session.
- An LU 6.2 application program, which provides a LOGON or SCIP exit routine to accept LU 6.2 sessions, can also use a confidential test option. You can code the CONFTXT=YES operand on the APPCCMD macroinstruction used to accept an LU 6.2 session (APPCCMD CONTROL=OPRCNTL, QUALIFY=ACTSESS). If you code CONFTXT=YES, VTAM clears its buffers that have been used to hold application program data before returning them to VTAM buffer pools for reuse.

Note: TSO also provides a CONFTXT option that you can use to protect confidential data. For information about CONFTXT and TSO, see [“Security” on page 567](#).

3270 Intrusion Detection Services

You can configure and monitor the VTAM 3270 Intrusion Detection Services (IDS) to determine problems in application 3270 data streams. The specific problem that is detected is the modification of protected fields in the data stream that 3270 emulators return.

This topic includes the following information:

- [“3270 IDS overview” on page 299](#) introduces the overview and background of the VTAM 3270 IDS function.
- [“3270 IDS considerations and assessment” on page 301](#) describes the various factors to consider before deploying the 3270 IDS solutions, including assessing your environment, deployment strategy, and known application 3270 solutions provided by IBM.
- [“Configuring 3270 IDS” on page 307](#) describes how to configure the monitoring of the 3270 data streams.
- [“Displaying and modifying 3270 IDS configuration” on page 309](#) describes the configuration and status of 3270 IDS.
- [“3270 IDS incidents” on page 311](#) describes the messages when 3270 data stream errors are detected.
- [“GTF trace data” on page 313](#) describes how to capture and format 3270 IDS incident trace data when 3270 IDS incidents are written to the GTF.
- [“Using SMF” on page 315](#) describes how to capture 3270 IDS incident trace data when 3270 IDS incidents are written to the SMF.
- [“Incident validation” on page 315](#) describes methods for gathering information about incidents.

3270 IDS overview

The z/OS Communications Server VTAM 3270 Intrusion Detection Services (IDS) function can help alert you to 3270 protocol violations as they occur in real time. This can be useful in identifying potential intrusions that attempt to manipulate 3270 protocol flows with the goal of compromising 3270 SNA applications and data that are deployed on your z/OS systems. This function can detect, in real time, an attempt by a malicious 3270 client emulator to modify protected fields on a 3270 screen. By modifying protected fields, the malicious 3270 client emulator might be trying to subvert the normal processing of the 3270 server application. The effect of such an attempt depends on how well the application guards

itself against unexpected changes to protected fields. In the best case scenario, a modification to a protected part of the screen is ignored by the application. In the worst case scenario, it could cause a potentially harmful change in the application's behavior.

Well behaved 3270 client emulator software typically prevents users from entering input into protected parts of the screen. The concern is over malicious users that use 3270 client emulators that do not honor the 3270 protocol and allow changes to protected fields. The 3270 IDS function can detect these types of protocol violations. However, note that SNA 3270 protocol violations might occur without malicious intent. This might be the result of race conditions or lax adherence of the SNA 3270 protocol by software such as 3270 client software emulators, the TN3270 client, session managers, or other SNA based 3270 protocol software. These anomalies might even occur with a regular frequency in your environment and most often go unnoticed as they do not have an impact that is visible to administrators, applications, or users. In some cases, they might cause a temporary error condition on the 3270 client's screen that they can easily recover from. While the 3270 IDS function can flag all detected protocol violations, it cannot determine whether a protocol violation is a malicious attack or an inadvertent anomaly in the 3270 protocol. Additionally, it cannot provide any insight on how a server-side 3270 application deals with these protocol anomalies. In other words, it cannot detect whether an application is vulnerable to a 3270 protocol-based attack or not. The 3270 IDS function simply detects and notifies system administrators of the presence of protocol anomalies, which can be useful as an audit log of potentially suspicious events. In addition to notification, the 3270 IDS function can be configured to take action on the SNA session when a protocol violation is detected, such as terminating the session.




Figure 91. 3270 IDS protection overview

Note: The z/OS Communications Server VTAM 3270 IDS solution is one of several solutions that can provide detection and protection from malicious 3270 attacks.

Figure 91 on page 300 provides an overview of the following 3270 data stream protocol validation solutions:

CICS basic mapping support (BMS)

CICS provides 3270 IDS detection and protection for any applications that exploit CICS basic mapping support (BMS) interfaces to create and parse their 3270 screens. When this support is activated, CICS monitors the 3270 data streams to detect any attempted modifications to protected fields on the screen. CICS can then provide warnings (log and error message) or prevent the application from processing the data by abending the transaction. See the CICS product documentation through the IBM Documentation: <https://www.ibm.com/docs/en> for more information on the CICS BMS IDS solution.

IMS Message Formatting Service (MFS) support

Similar to the CICS BMS, IMS provides 3270 IDS support for any IMS applications that use the IMS Message Format Service (MFS) to format and parse their 3270 messages. When this function is enabled, IMS prevents modifications to protected fields from being passed on to IMS server applications. See the IMS product documentation through the IBM Documentation: <https://www.ibm.com/docs/en> for more information on the IMS MFS IDS solution.

VTAM 3270 IDS support

The VTAM 3270 IDS support is described in this topic.

ISPF also provides built-in IDS support. ISPF is one of the other subsystems shown in [Figure 91 on page 300](#). Applications that use ISPF services to display their 3270 panels are automatically protected by ISPF. ISPF automatically detects and prevents any modifications to protected areas of the panels from occurring.

The list of 3270 data stream protocol validation solutions is not intended to be an exhaustive list. The other subsystems or other middleware category shown in [Figure 91 on page 300](#) is intended to indicate any other potential application layer 3270 IDS support that might exist but is not identified here.

Note:

- The 3270 client emulators that are used by the 3270 users can use native SNA attachment directly to VTAM or IP attachment through TN3270. The VTAM and middleware 3270 IDS support that is shown in [Figure 91 on page 300](#) covers all 3270 users.
- The terminology in this topic refers to general 3270 validation support, which is different from the specific terminology, such as CICS BMS, IMS MFS, or VTAM 3270 IDS, which refers to validation support within specific products that support the 3270 protocol.

3270 IDS considerations and assessment

This topic describes the various factors that you should consider, steps for assessing your exposure to potential 3270 protocol-based attacks, and your potential need for deploying one or more of the 3270 IDS solutions that are described in [“3270 IDS overview” on page 299](#).

Assessing your environment

If you have workloads that are protected by middleware or native application 3270 validation, evaluate the solutions that are described in [“3270 IDS overview” on page 299](#) first before you investigate the z/OS Communications Server VTAM 3270 IDS function. Generally, the application solutions have a much lower overhead for IDS processing than the z/OS Communications Server VTAM solution, as they already have existing processing for the processing and handling the 3270 data streams.

The z/OS Communications Server VTAM 3270 IDS function can complement these solutions if you have workloads that you determine are at risk and are not covered by other IDS solutions. As a result, the z/OS Communications Server VTAM 3270 IDS solution is not necessarily required by all z/OS systems or users who have SNA 3270 application workloads. As with any intrusion detection capabilities, you must take careful considerations before you enable a 3270 IDS function. For this reason, the background information is provided here for you to analyze and determine whether this type of IDS function can provide value to your environment. As part of this analysis, you need to understand the z/OS Communications Server VTAM 3270 IDS function, other 3270 IDS options and applicability to your environment, and then assess the risk and cost factors in your environment.

This topic provides information to assist you in your assessment for the potential need of this function in your z/OS environment by evaluating the following aspects:

Applications

Identify candidate applications and perform a careful analysis of need. See [SNA application applicability criteria](#).

End users, SNA technology and connectivity

Evaluate the users, key SNA technologies, and network configuration considerations. See exploitation factors and their implications in [“SNA technologies, network connectivity, and environmental factors”](#) on page 304.

Exploitation cost

Understand the cost to exploit this function. See system resource cost and administrative considerations in [Exploitation cost](#).

Complete this assessment carefully, and then consider moving forward with the exploitation of the 3270 IDS function.

SNA application applicability criteria

Is the 3270 IDS function needed in your environment?

Many factors help determine the need for the validation protection that is provided by the 3270 IDS function. To make this assessment, you need background in securing SNA workloads. See the [3270 Emulation: Security Considerations](#) white paper for initial information.

After reviewing the security information in this white paper, you can continue your assessment by using the following key 3270 IDS considerations.

Carefully evaluate your SNA applications for each z/OS system. This topic provides considerations for your SNA application workloads (per z/OS system).

SNA applications

All z/OS systems have 3270 application workloads. At a minimum, the system administrators use various TSO/ISPF functions to maintain z/OS and often to manage other applications. Beyond the system administrators, various applications can exploit SNA APIs that are related to SNA LU0 and LU2 3270 workloads. The first step in your assessment is to identify all of the 3270 applications on your applicable z/OS systems. To assist with this step, use the VTAM operator display command **D NET,APPLS,SCOPE=3270CAND**. This display provides the following information:

3270 candidate applications

Displays a list of active VTAM applications that have any LU-LU sessions (since the ACB was opened) that qualify for the 3270 IDS monitoring. The qualifying LU sessions must be LU type LU0 with TS profile 2 or LU2 with TS profile 3. Applications that have qualifying LU sessions are potential candidates for the 3270 IDS function.

LU session count

A cumulative session count (since the ACB was opened) of the number of qualifying LU sessions.

While all applications in this list are candidates, you should initially focus on the applications with the highest qualified LU session counts. After you identify your candidate 3270 applications, you need to evaluate the application 3270 support for each of those applications.

After you identify the applications to focus on, consider whether the VTAM 3270 applications themselves or the middleware under which they run, for example, IBM middleware such as CICS or IMS, offer any native 3270 protocol related validation or protection. As described in [“3270 IDS overview”](#) on page 299, CICS and IMS provide modes (BMS and MFS) for their applications to exploit 3270 communications that also provide 3270 protocol validation. You should first evaluate the data validation support that is provided by the IBM middleware and possibly by the application programs themselves. You might need to consult with the CICS or IMS application programming staff to understand what modes are exploited. If the VTAM 3270 application does not run under a middleware environment that provides its own 3270 IDS function, you need to consult with the application support staff or supporting documentation for the application.

If protocol validation is offered by the middleware or application, enable its support. The 3270 middleware or application is typically in a better position to perform this type of protocol validation. With an existing understanding of the 3270 data stream context, middleware and application validation is typically much more efficient than the VTAM IDS approach. The application can also provide for some error recovery, retries, or have the capability to ignore certain anomalies or error conditions.

Figure 92 on page 303 provides an overview of the candidate application assessment process. For each candidate 3270 application, start your assessment here.

Start your assessment here




Figure 92. Candidate application assessment process

After you complete your assessment of the candidate applications and their native 3270 validation support, you might determine that the VTAM 3270 IDS validation is not required for your list of identified applications. If you do have a list of candidate applications without coverage or you have a list of potential candidate applications that have unknown validation capability, continue your assessment.

3270 user community of end users

For the identified candidate applications, how well do you understand the configuration and access of your 3270 emulators and the actual end users? For example, who are your end users for each application? Are the users internal or within your company, or are some of them external users? How many users are there?

For the users of each application, what forms of access control and authentication do you have in place for the 3270 users, for example, TLS/SSL, SAF-based, custom written, and so on?

More SNA related aspects are also end user considerations:

- Can you identify or inventory the various SNA components that are used for host access by this set of users?
- What is your level of control and confidence for the security of the following components:
 - TN3270 server products
 - TN3270 client products
 - What products do the SNA session managers support?
 - Do the SNA native connections use TN3270 access? If yes, what 3270 access solutions are being used?

You might not need the 3270 validation for this application, if both the following conditions are met:

- The scope of your end users who have access to your 3270 applications is known and limited.
- The level of control or trust (authentication) you have with the access for those end users (including the control over and integrity of the TN3270 client software and protection it provides) is high.

Consider performing a risk assessment to determine whether the additional protection of an IDS solution is warranted for this application by considering this set of users and the client software used by the users. To complete this part of the assessment, you might need to consult with the 3270 client emulator vendor.

See [“SNA technologies, network connectivity, and environmental factors” on page 304](#) for more end user related considerations.

SNA technologies, network connectivity, and environmental factors

[Figure 93 on page 305](#) illustrates a typical SNA 3270 network configuration that shows an SNA session manager and a TN3270 server that are located within the same z/OS instance. Several key environmental and configuration aspects are related to identifying and reporting 3270 protocol violations. Some aspects overlap with the previous end user considerations.




Figure 93. Sample of typical SNA 3270 network configuration

Many variables in an SNA 3270 network can impact the flow of the 3270 data stream. Among these are layers of SNA components, connectivity, and often SNA session management products. The session flow can potentially have an impact on the 3270 protocol validation processing. Consider your unique environmental aspects, for example:

- Systems or network topology: CPU utilization/availability and network topology, network configuration, and network equipment, which can all impact latency that impacts timing.
- Your 3270 related products and vendors, for example, the 3270 application (IMS, CICS, and so on), the TN3270 servers, whether on z/OS or other platforms, the 3270 clients, possible SNA session managers, your end users, and native SNA connectivity when applicable.
- Product compliance, level and compatibility of the SNA LU0 and LU2 3270 protocols, each product in the path of the 3270 data stream, including the 3270 client emulator product, adherence to and implementation of the SNA LU0 and SNA LU2 3270 architecture (including any applicable SNA extensions).
- Your 3270 system and network configuration, physical proximity of resources and IP topology (relative to the 3270 application), the location of the session manager, the TN3270 server and platform, the 3270 client (distance and other network topology aspects) that can all impact timing.

If you have SNA session managers that run on your z/OS system as illustrated in Figure 93 on page 305, you should disable VTAM 3270 IDS for the traffic between the 3270 clients and the session managers to avoid double validation for the same 3270 data streams. Instead, only perform the validation between the session manager and the actual 3270 application.

It is important to understand that the VTAM 3270 IDS function reports any violation of the 3270 protocol that causes a protected field to be overwritten. The violation might be the result of malicious activity or the result of unintentional protocol anomalies (inadvertent or transient protocol violations that are caused by things like timing or queuing anomalies).

The VTAM 3270 IDS function cannot make the distinction between a malicious activity versus an unintentional protocol anomaly. Instead, such distinction requires careful analysis of the captured documentation that is associated with the reported incident. In many cases, this type of error can be handled (ignored or retried) by the application.

The frequency of reporting unintentional protocol anomalies varies for each environment. In some environments, the amount or frequency of reporting of unintentional protocol anomalies can be problematic. For each reported incident, you need to perform the initial evaluation to determine the disposition of the IDS incident. If the reported incident is a known unintentional protocol anomaly, that type of incident can be self-managed. If the reported incident is determined to be a malicious attack, you can use the information that is recorded by the VTAM 3270 IDS function to begin your effort to identify the source of the attack. Finally, you might determine that the reported incident is a valid use of the 3270 protocol even though the VTAM 3270 IDS support flagged it. In that case, you might need to work with IBM service to determine the disposition of the incident.

Exploitation cost

System resource cost and other implications of exploitation

Consider system resources (CPU and storage) and administrative costs that are related to the exploitation of the VTAM 3270 IDS function. The actual performance impact of enabling the function varies for each customer environment depending on the scope of the support enabled the application workloads and the type of 3270 traffic.

Processing cost

Internal IBM benchmarks indicate that the IDS analysis that is performed by VTAM 3270 IDS function can result in an increase in CPU use for the SNA application address space, for example, the CICS TOR address space. The amount of increase is impacted by several factors such as the format, complexity, and size of the 3270 screens typically used by the application. The number of LU sessions is another key factor.

If you have SNA session managers that run on your z/OS system, you should disable VTAM 3270 IDS for the traffic between the 3270 clients and the session managers to avoid double validation for the same 3270 data streams. Instead, only perform the validation between the session manager and the actual 3270 application.

Virtual memory cost

Each SNA session that is enabled for the 3270 IDS function allocates approximately 100 K for the session and extra storage for outbound PIU tracking. The VTAM DSCOUNT start option determines how many outbound PIUs to save. The DSCOUNT setting along with PIU (screen) size directly affects the amount of virtual memory that is used to monitor the session. The session-related storage is all 64-bit virtual storage. Total virtual storage can be estimated by multiplying 125 K by total sessions monitored. You need to insure that enough real and virtual memory (paging space) is available before you enable this function on a system. Additionally, insure that system parameters such as MEMLIMIT are set to appropriate values that do not artificially limit the amount of 64-bit storage that is available to the relevant address spaces.

Administrative cost

Administrative costs are associated with your initial applicability analysis, enablement, and monitoring. Some coordination might be required with the applications administrative staff as well. After the VTAM 3270 IDS function is enabled, each reported incident provides diagnostic data that needs to be evaluated by your staff. This might include network administrators, application developers, other personnel who are familiar with the 3270 data streams in question. If your evaluation concludes that the incident does not reflect any type of protocol violation, you can contact IBM for further assistance.

In some cases, the exploitation of this IDS function might result in persistent and ongoing reporting of similar unintentional protocol anomalies; for example, due to specific vendor products such as TN3270 server, client emulation support, session managers. In such cases, you are required to work with the associated vendor product support staff for a resolution. Pending a resolution, the VTAM 3270 IDS function can be disabled for such workloads.

IBM provides changes for the z/OS TN3270 Telnet Server and the CS Distributed Telnet Server that are related to timing scenarios that can result in reporting of protocol anomalies. For more information about those issues and related changes, see the product support information for those products.

Deployment strategy

After you complete the assessment for each z/OS system and the applicable application workloads on those systems and you believe that your environment can benefit from the VTAM 3270 IDS function, you should consider creating an exploitation plan that enables the support on your systems in a controlled and gradual manner in terms of systems and applications. The objective is to gradually extend the support to the various 3270 application workloads. For example, the support should start with test or development systems for specific applications. After the support is active for a period of time and you have assessed the impact, you can continue expanding the support to other workloads and systems.

Do not code the DSACTION=SENSE option that raises error condition to the z/OS 3270 application or the DSACTION=TERM option that terminates corresponding LU-LU session until you have sufficient experience with a workload and justification for this setting. While you expose more systems and workloads to this validation support, you can assess the impact to your environment and the effectiveness for your workloads.

Known application 3270 solutions provided by IBM

CICS and IMS application 3270 validation support

Both IMS and CICS products provide the following native 3270 validation support:

- CICS BMS data stream protection
- IMS MFS data stream protection

Both of these solutions require less processing and CPU than the VTAM 3270 IDS solution. For more information about the CICS or IMS support, see the CICS or IMS product documentation for related PTF or base product information.

Application API dynamic control

The VTAM IDS support also provides the capability for applications to dynamically control (enable/disable) the VTAM IDS validation function. With this support, middleware-based IDS solutions can temporarily disable the VTAM IDS function when the middleware IDS solution is actively protecting a session or avoiding dual monitoring of the session when both middleware and VTAM IDS solutions are active for a session. IBM middleware products, such as CICS, IMS, and ISPF, provide this support. For more information about the dynamic IDS exploitation provided by CICS or IMS, see the related product documentation.

TSO/ISPF considerations

TSO users, who are in the ISPF environment, are protected by the ISPF built-in 3270 validation support that is always active. Other TSO environments (non ISPF) need to be investigated and can be a candidate for the z/OS Communications Server VTAM 3270 IDS support. ISPF uses the application API dynamic control when processing ISPF panels.

Configuring 3270 IDS

About this task

To control the monitoring of 3270 data streams, you can use the following start options or specify the following parameters on the GROUP and APPL statements in an application major node.

DSMONITR

Controls the basic function of the monitoring.

DSACTION

Controls the actions that will be taken when an incident occurs.

DSCOUNT

Controls the number of outbound 3270 data streams to help reconstruct the events that cause the incident.

DSTRUST

Controls the type of secondary logical units that are trusted and therefore not monitored.

Procedure

Sample configuration on the ATCSTRDS, ATCCONDS, and TSO1A parameters shows the definitions that are required for the 3270 IDS monitor function. This sample configuration enables the following functions:

- Only the TSO applications are monitored.
- The local logical units are trusted, but others are not.
- Message group IST2424I is written to the console.
- No additional action will be taken when an incident is detected.
- If more than 10 incidents occur within 60 seconds, writing the message group IST2424I stops for the rest of those 60 seconds.
- Up to 15 outbound buffers are saved for each session that is monitored.

Perform the following steps to configure 3270 Intrusion Detection Services:

1. Configure the ATCSTRDS member with the DSMONITR keywords.

```
*****
*
*  SAMPLE PARAMETER FOR THE 3270 INTRUCTION DETECTION SERVICES
*
*****
DSMONITR=APPL,          MONITOR ONLY APPLS WITH DSMONITR=YES
DSACTION=(CONSOLE,NONE,10),  MESSAGES TO THE CONSOLE
DSCOUNT=15,              SAVE UP TO 15 OUTBOUND BUFFERS
DSTRUST=NONE,            COLLECT FOR ALL TYPES OF SESSIONS
CONFIG=DS,              START CONFIG
...
Rest of the VTAM start member
```

2. Configure the VTAM member ATCCONDS to define the TSO application major node at start up.

```
*****
*
*  NAME:  ATCCONDS
*
*  USE:   CONFIGURATION LIST FOR SSCP1A.
*
*
*  SECURITY: IBM INTERNAL USE ONLY
*****
          TSO1A,              TSO Applications
...
X
```

3. Configure the VTAM TSO application major node.

```

*****
*
* NAME:  TS01A
*
* USE:   APPL DECK FOR TSO
*
*****
      VBUILD TYPE=APPL
      GROUP DSMONITR=YES,DSTRUST=LOCALLU
*****
* THE FOLLOWING APPLS ARE FOR TSO/VTAM
*****
TS01      APPL  ACBNAME=TSO,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1,      C
          DSMONITR=NO
TS010001 APPL  ACBNAME=TS000001,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010002 APPL  ACBNAME=TS000002,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010003 APPL  ACBNAME=TS000003,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010004 APPL  ACBNAME=TS000004,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010005 APPL  ACBNAME=TS000005,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010006 APPL  ACBNAME=TS000006,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010007 APPL  ACBNAME=TS000007,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010008 APPL  ACBNAME=TS000008,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010009 APPL  ACBNAME=TS000009,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
TS010010 APPL  ACBNAME=TS000010,AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1

```

In this sample definition, the GROUP DSMONITR value is set to YES, but the APPL DSMONITR definition for TSO1 is set to NO. Because TSO does a CLSDST PASS to the other TSO applications, no 3270 monitoring is needed for this application.

Displaying and modifying 3270 IDS configuration

You can use VTAM commands to display the status of 3270 monitoring or to change the 3270 monitoring configuration.

- Start VTAM. No messages are issued for the 3270 IDS function during VTAM start up.

```

S NET,,, (LIST=DS)
IEF403I NET - STARTED - TIME=09.58.50
IST1054I VALUE FOR SIZE MUST BE BETWEEN 4M AND 2048M
IST448I DSPSIZE OPTION IGNORED - NO LONGER SUPPORTED
IST093I ISTCDRDY ACTIVE
IST315I VTAM INTERNAL TRACE ACTIVE - MODE = INT, SIZE = 0004 MB
IST199I OPTIONS = API APPC CFS CIA CIO CSM ESC HPR LOCK MSG
IST199I OPTIONS = NRM PIU PSS SMS SSCP TCP VCNS XBUF XCF
IST314I END
...
IST093I TS01A ACTIVE
...
IST020I VTAM INITIALIZATION COMPLETE FOR CS 3.1
IST1349I COMPONENT ID IS 5695-11701-310
IST1348I VTAM STARTED AS INTERCHANGE NODE

```

- Issue the **Display NET,VTAMOPTS,FUNCTION=SECURITY** command to display the current values of the DSMONITR keywords.

```

DISPLAY NET,VTAMOPTS,FUNCTION=SECURITY
IST097I DISPLAY ACCEPTED
IST1188I VTAM CS 3.1 STARTED AT 09:58:50 ON 01/25/16 655
IST1349I COMPONENT ID IS 5695-11701-310
IST1348I VTAM STARTED AS INTERCHANGE NODE
IST1189I DSACTION = (CONSOLE,NONE,10)      DSCOUNT   = 15
IST1189I DSMONITR = APPL                    DSTRUST   = NONE
IST1189I ENCRPREF = NONE                     ENCRYPTN   = 31
IST1189I IPINFO   = SENDALL                  SECLVLCP  = ***NA***
IST1189I VERIFYCP = NONE
IST314I END

```

- Issue the **MODIFY proc,DSMONITR=NO** command to turn off the DSMONITR function. The **MODIFY** command stops the 3270 IDS monitoring. When secondary logical units send or receive a 3270 data stream, monitoring for the session is stopped. If the value of the DSMONITR keyword is changed back to YES, monitoring starts when a secondary logical unit starts a new session with the application.

```

MODIFY NET,VTAMOPTS,DSMONITR=NO
IST097I MODIFY ACCEPTED
IST223I MODIFY COMMAND COMPLETED

```

- Issue the **Display VTAMOPTS** command to display the updated values.

```

D NET,VTAMOPTS,FUNCTION=SECURITY
IST097I DISPLAY ACCEPTED
IST1188I VTAM CS 3.1 STARTED AT 09:58:50 ON 01/25/16 946
IST1349I COMPONENT ID IS 5695-11701-310
IST1348I VTAM STARTED AS INTERCHANGE NODE
IST1189I DSACTION = (CONSOLE,NONE,10) DSCOUNT = 15
IST1189I DSMONITR = NO DSTRUST = NONE
IST1189I ENCRPREF = NONE ENCRYPTN = 31
IST1189I IPINFO = SENDALL SECLVLCP = ***NA***
IST1189I VERIFYCP = NONE
IST314I END

```

- Issue the **Display NET,ID=applname** command to display the status of an application. The status of ACTIV/3-S indicates that this session is being monitored.

```

D NET,ID=TS00002,SCOPE=ALL
IST097I DISPLAY ACCEPTED
IST075I NAME = TS00002, TYPE = APPL 986
IST486I STATUS= ACT/S, DESIRED STATE= ACTIV
IST1447I REGISTRATION TYPE = NO
IST599I REAL NAME = NETA.TS0100021
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST1632I VPACING = 7
IST1938I APPC = NO
IST597I CAPABILITY-PLU ENABLED ,SLU DISABLED ,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = TS01A
IST213I ACBNAME FOR ID = TS010001
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = JHACKER, STEPNAME = OS390R5, DSPNAME = ISTFF999
IST228I ENCRYPTION = OPTIONAL , TYPE = DES
IST1563I CKEYNAME = TS010001 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = 0 MAXIMUM = 0
IST2433I DSMONITR = YES, DSCOUNT = 15, DSACTION = (CONSOLE,NONE)
IST2434I DSTRUST = LOCALLU
IST2435I SESSIONS MONITORED = 1, ERRORS DETECTED = 0
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME STATUS SID SEND RECV VR TP NETID
IST635I TCPM0001 ACTIV/3-S EAABEEC331E8DB02 0016 0003 NETA
IST314I END

```

- Issue the **D NET,SESSIONS** command to display information about the session. Message IST2436I or IST2437I shows the status of the 3270 IDS monitor.

```

D NET,SESSIONS,SID= EAABEEC331E8DB02
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 056
IST879I PLU/DLU REAL = NETA.TS010001 ALIAS = NETA.TS00002
IST879I SLU/OLU REAL = NETA.L7201A ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV/3
IST933I LOGMODE=D4B32XX3, COS=BLANK*
IST1635I PLU HSCB TYPE: FMCB LOCATED AT ADDRESS X'176F5188'
IST1635I SLU HSCB TYPE: LUST LOCATED AT ADDRESS X'175A9314'
IST2437I DSMONITR = YES, ERRORS DETECTED = 0
IST2064I PLU TO SLU RU SIZE = 65535 SLU TO PLU RU SIZE = 6144
IST1636I PACING STAGE(S) AND VALUES:
IST1637I PLU--STAGE 1--SLU
IST1638I STAGE1: PRIMARY TO SECONDARY DIRECTION - FIXED
IST1639I PRIMARY SEND: CURRENT = 0 NEXT = 0
IST1640I SECONDARY RECEIVE = 0
IST1641I STAGE1: SECONDARY TO PRIMARY DIRECTION - NO PACING
IST1714I NO PATH INFORMATION EXISTS
IST314I END

```

- Issue the **D NET,STATS,TYPE=VTAM** command to display statistics about IDS activity.

```

D NET,STATS,TYPE=VTAM
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = STATS,TYPE=VTAM 308
IST1349I COMPONENT ID IS 5695-11701-310
IST1345I ID VALUE DESCRIPTION
IST1227I 151 0 = DEPENDENT LU TOTAL FOR ISTPUS
IST1227I 11 0 = CHANNEL-TO-CHANNEL ATTACHMENTS
IST1227I 61 0 = SNA DATA COMPRESSION SESSIONS
IST1227I 63 24 = RECOVERABLE SESSIONS
...
IST1227I 170 1 = IDS3270 TOTAL SESSIONS MONITORED
IST1227I 171 1 = IDS3270 CURRENT SESSIONS MONITORED
IST1227I 172 1 = IDS3270 SESSIONS MONITORED SINCE ENABLE
IST1227I 173 0 = IDS3270 TOTAL INCIDENTS FOUND
IST1227I 174 0 = IDS3270 TOTAL SUPPRESSED CONSOLE REPORTS
IST1454I 91 STATISTICS DISPLAYED

```

- Display the status of CSM storage.

Tip: If too much HVCOMM storage is in use, consider reducing the DSCOUNT value.

```

D NET,CSM
IVT5508I DISPLAY ACCEPTED
IVT5529I PROCESSING DISPLAY CSM COMMAND - OWNERID NOT SPECIFIED
IVT5530I BUFFER BUFFER
IVT5531I SIZE SOURCE INUSE FREE TOTAL
IVT5532I -----
...
IVT5532I -----
IVT5533I 4K HVCOMM 4K 1020K 1M
IVT5533I 16K HVCOMM 0M 0M 0M
IVT5533I 32K HVCOMM 32K 992K 1M
IVT5533I 60K HVCOMM 120K 900K 1020K
IVT5533I 180K HVCOMM 0M 0M 0M
IVT5535I TOTAL HVCOMM 156K 2912K 3068K
...
IVT5604I HVCOMM MAXIMUM = 2000M HVCOMM CURRENT = 3M
IVT5541I HVCOMM MAXIMUM USED = 3M SINCE LAST DISPLAY CSM
IVT5594I HVCOMM MAXIMUM USED = 3M SINCE IPL
...
IVT5599I END

```

3270 IDS incidents

If the 3270 IDS monitor detects a 3270 data stream error, an incident report is issued.

```

IST2424I 3270 DATA STREAM ERROR - NETA.TS00002 NETA.TCPM0001
IST2425I PLU SUBAREA = X'0001' INDEX = X'0001' ELEMENT = X'0076'
IST2425I SLU SUBAREA = X'0001' INDEX = X'0000' ELEMENT = X'003A'
IST2441I JOBNAME = JHACKER SID = EAABEEC331E8DB02
IST2426I IPADDR = 192.168.98.254..61691
IST2427I DATE = 2016/01/25 TIME = 15:47:56 ID = 1
IST2428I ROW = 9 COLUMN = 16
IST2429I OUTBOUND - SEQ = X'0001' OFF = 598 LEN = 39
IST2431I 40404040 40404040 D1C1C3D2 E2D6D540 * JACKSON *
IST2430I INBOUND - SEQ = X'0001' OFF = 284 LEN = 39
IST2431I 40404040 40404040 F1F2F3F4 F5F6F7F8 * 12345678*
IST314I END

```

Messages IST2424I, IST2425I, IST2441I, IST2426I, and IST2427I describe the identifying information about the incident. Messages IST2428I, IST2429I, IST2430I, and IST2431I describe the information about the specific data that caused the incident. This same information is written to the Generalized Trace Facility (GTF) by using the Event ID (EID) value x'F90'.

Use automation to start GTF when message IST2430I is issued in case any additional incidents exist. It is recommended that you have GTF running before the message IST2430I is issued to avoid the need for a recreate. In addition, VTAM buffer trace can be started for the secondary logical unit.

Tip: If the DSACTION=SYSLOG option is active, an IST2424I message is written to the console and the IST2424I message group is written to SYSLOG. The IST2424I message appears twice for the same incident in the automation program.

If the number of incidents that occur in a 60-second interval is more than the *msg-count* parameter of the DSACTION start option, the IST2424I message groups are not displayed on the console. When the interval expires, the message group IST2432I is displayed on the console.

```
IST2432I 3270 ERROR SUMMARY FROM 2016/01/25 AT 15:49:34
IST2438I SESSIONS = 1 ERRORS = 3
IST924I -----
IST2439I PLU          SLU          SID          ERRORS
IST2440I NETA.TS010002  NETA.TCPM0001  EAABEEC331E8DB02  3
IST314I END
```

Use the **CSDUMP** or **SLIP** command to capture a dump when an incident occurs.

Example when you use the **CSDUMP** command.

```
F NET,CSDUMP,MESSAGE=IST2430I
IST097I MODIFY ACCEPTED
IST223I MODIFY CSDUMP COMMAND COMPLETED
D NET,CSDUMP
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = CSDUMP TRIGGERS 993
IST1871I MESSAGE TRIGGER: MESSAGE = IST2430I MATCHLIM = 1
IST1875I SENSE TRIGGER: NONE
IST314I END

...
IEA045I AN SVC DUMP HAS STARTED AT TIME=13.37.09 DATE=02/08/2016 046
FOR ASIDS(0031,001D)
QUIESCE = YES
IST1879I VTAM DUMPING FOR CSDUMP TRIGGER MESSAGE IST2430I
IST2430I 3270 DATA STREAM ERROR - NETA.TS010002 NETA.TCPM0001 048

...
IST314I END
IEA794I SVC DUMP HAS CAPTURED: 049
DUMPID=002 REQUESTED BY JOB (USER2 )
DUMP TITLE=ISTRACSW - MSG CSDUMP - ID=6C5C
IEA911E COMPLETE DUMP ON SYS1.DUMP01 055
DUMPID=002 REQUESTED BY JOB (USER2 )
FOR ASIDS(0031,001D)
INCIDENT TOKEN: LOCAL VIC000 02/08/2016 18:37:09
```

Example when you use the **SLIP** command.

```
SLIP SET,ENABLE,MSGID=IST2430I,ID=2430,ACTION=SVCD,JOBNAME=NET,
END
IEE725I SLIP PARAMETERS ARE- 142
ID=2430,NONPER,ENABLED
ACTION=SVCD,SET BY CONS IC000A,RBLEVEL=ERROR,MATCHLIM=1,0
JOBNAME=NET,MSGID=IST2430I
IEE727I SLIP TRAP ID=2430 SET

...
IEA045I AN SVC DUMP HAS STARTED AT TIME=14.06.53 DATE=02/08/2016 165
FOR ASID (001D)
QUIESCE = YES
IEA992I SLIP TRAP ID=2430 MATCHED. JOBNAME=NET , ASID=001D.
IEA411I SLIP TRAP ID=2430 DISABLED FOR MATCHLIM
IEA794I SVC DUMP HAS CAPTURED: 168
DUMPID=003 REQUESTED BY JOB (NET )
DUMP TITLE=SLIP DUMP ID=2430
IST2424I 3270 DATA STREAM ERROR - NETA.TS010002 NETA.TCPM0001 164

...
IST314I
END
```

The VTAMMAP SES formatted dump tool will format the 3270 incidents that are found for each session. See [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#) for more information.

GTF trace data

When a 3270 IDS incident occurs, trace records are written to an active generalized trace facility (GTF).

Use the following command to start GTF. To prevent prompting, add the NOPROMPT option to the GTF procedure parameter, which enables starting an automated GTF procedure.

```
S GTF.GTF,DSN=USER.TRACE,DISP=OLD,MEMBER=GTFF90
AHL121I TRACE OPTION INPUT INDICATED FROM MEMBER GTFF90 OF PDS SYS1.PARMLIB
AHL103I TRACE OPTIONS SELECTED--USR
00 AHL125A RESPECIFY TRACE OPTIONS OR REPLY U
REPLY 00,U
AHL031I GTF INITIALIZATION COMPLETE
```

Use the following command to stop GTF.

```
P GTF
AHL006I GTF ACKNOWLEDGES STOP COMMAND
AHL904I THE FOLLOWING TRACE DATASETS CONTAIN TRACE DATA :
      USER.TRACE
```

For more information about using GTF, see *The Generalized Trace Facility (GTF)* in [z/OS MVS Diagnosis: Tools and Service Aids](#).

Collecting GTF trace data

Event ID (EID) records are always written to available GTF which allows the writing of the EID when a 3270 IDS incident occurs. Up to DSCOUNT outbound PIUs and the inbound PIU that caused the incident are written.

The following example shows a sample GTF procedure.

```
//GTFNEW PROC MEMBER=GTFFPARM
//IEFPROC EXEC PGM=AHLGTF,'PARM=MODE=EXT,DEBUG=NO,TIME=YES',
// TIME=1440,REGION=4M
//IEFRDRE DD DSN=SYS1.TRACE,UNIT=SYSDA,SPACE=(TRK,20),
// DISP=(NEW,KEEP)
//SYSLIB DD DSN=SYS1.PARMLIB(&MEMBER),DISP=SHR
```

The following example shows a sample GTF SYS1.PARMLIB(GTFF90) member. F90 is the EID of the 3270 IDS trace records. FEF, FF0, and FF1 are VTAM buffer trace EIDs.

```
TRACE=USRP
USR=(F90,FEF,FF0,FF1)
```

Formatting GTF trace data

Use the **IPCS GTFTRACE** command to format the collected generalized trace facility (GTF) data for a 3270 Intrusion Detection Services (IDS) incident.

```
GTFTRACE DSN('USER.TRACE') USR(F90)
```

For more information about the GTFTRACE command, see [z/OS MVS IPCS Commands](#).

For each 3270 IDS incident, up to DSCOUNT outbound PIUs are traced. The inbound PIU, which contained the data stream that caused the incident to be found and recorded, is also traced. Each trace record contains information about the incident.

The following example shows the formatted 3270 IDS trace records.

```

USRFD F90 ASCB 00F8EE00          JOBN JHACKER
                                     **** 3270 Data Stream Error ****
(1)3270   NETA.TCPM0001   /NETA.TS00002   LRC(000,000)   OUTBOUND   COMPLETE SEGMENT
(2)Time   UTC 2016/01/25 20:47:56.476213   LOC 2016/01/25 15:47:56.476213
(3)Event   Token 0000000001   SID EAABEEC3   31E8DB02   Buffer 01 of 01
(4)IPAddr  192.168.98.254..61691
(5)Overlap Row 009 Col 016 Offset 00665
(6)OUT     SEQ X'0001'   Offset 00598   Length 00039
(7)         40404040 40404040 D1C1C3D2 E2D6D540 40404040 * JACKSON *
         40404040 00000000 00000000 * ..... *
(8)IN      SEQ X'0001'   Offset 00284   Length 00039
(9)         40404040 40404040 F1F2F3F4 F5F6F7F8 F9404040 * 123456789 *
         40404040 114AE9F6 F14040D7 * .Z61 P *
(10)Buffer UTC 2016/01/25 20:47:26.450328   LOC 2016/01/25 15:47:26.450328
(11)VTAM   TH=40000000 00000000 00010001 00000001 1800000B 00580001 051F RH=0380C0
(12)       SEQ 0001-0001   F5C21140 402901C0 40F4F040 40E44040 40404040 *5B. ...{ 40 U *
         404040C3 C8D9C9E2 E3C9C1D5 40404040 40404008 * CHRISTIAN *
...
         114DC829 01C0E9C5 F94040D7 40C8E240 40D44040 *.H..{ZE9 P HS M *
         40D4C1E2 D6D54040 40404040 40404011 4DF02901 * MASON *(0..*
         C06CF6C3 4040D740 4040C940 40404040 D1C1C3D2 *}%6C P I JACK*
         E2D6D540 40404040 40404040 114ED829 01C06DF6 *SON .+Q..{6*
...
         40404040 40404040 40C8C5E7 E2E3D9C9 D5C74DF0 * HEXSTRING(0*
         F05D4011 5D7E1D60 *0) .)=.- *
(13)      GMT-01/25/2016 20:47:56.476251   LOC-01/25/2016 15:47:56.476251

USRFD F90 ASCB 00F8EE00          JOBN JHACKER
                                     **** 3270 Data Stream Error ****
(1)3270   NETA.TS00002   /NETA.TCPM0001   LRC(000,000)   INBOUND   COMPLETE SEGMENT
(2)Time   UTC 2016/01/25 20:47:56.476213   LOC 2016/01/25 15:47:56.476213
(3)Event   Token 0000000001   SID EAABEEC3   31E8DB02   CODE U('E4')
(4)IPAddr  192.168.98.254..61691
(5)Overlap Row 009 Col 016 Offset 00665
(6)OUT     SEQ X'0001'   Offset 00598   Length 00039
(7)         40404040 40404040 D1C1C3D2 E2D6D540 40404040 * JACKSON *
         40404040 00000000 00000000 * ..... *
(8)IN      SEQ X'0007'   Offset 39044   Length 00001
(9)         40404040 40404040 F1F2F3F4 F5F6F7F8 F9404040 * 123456789 *
         40404040 114AE9F6 F14040D7 * .Z61 P *
(10)Buffer UTC 2016/01/25 20:47:56.476216   LOC 2016/01/25 15:47:56.476216
(11)VTAM   TH=40000000 00000000 00000001 00010001 1C000058 000B0001 0298 RH=0393A0
(12)       SEQ 0001-0001   7D4AD811 40E9C3F1 4040E440 40404040 D4404040 *'Q. ZC1 U M *
         C1D3C5E7 E8E24040 40404040 40404040 11C1F9C3 *ALEXYS .A9C*
         F54040E4 4040E240 40D44040 40D4C1E2 D6D54040 *5 U S M MASON *
         40404040 40404040 4011C3C9 C3F94040 E440C8E2 * .CIC9 U HS*
...
         40E4D540 40C940D4 404040D4 C1E2D6D5 40404040 * UN I M MASON *
         40404040 40404011 4AC1F6F0 4040D740 40404040 * .A60 P *
         40404040 F1F2F3F4 F5F6F7F8 F9404040 40404040 * 123456789 *
         114AE9F6 F14040D7 40404040 40D44040 40D4C1C4 *.Z61 P M MAD*
         C9E2D6D5 40404040 40404040 401148F9 C5F54040 *ISON ..9E5 *
...

C8C540E5 C1D3E4C5 40E3D67A 40404040 40404040 *HE VALUE TO: *
40404040 404040C8 C5E7E2E3 D9C9D5C7 4DF0F05D * HEXSTRING(00)*
40 * *
(13)      GMT-01/25/2016 20:47:56.476251   LOC-01/25/2016 15:47:56.476251

```

In the example:

- (1) The network names of the primary logical unit (PLU) and secondary logical unit (SLU), the lost record counts, the direction of the packet (inbound or outbound), and the position of the RU in the traced records. Outbound packets trace the entire chain of RUs from the begin chain to the end chain. Inbound packets trace only the specific RU that caused the incident.
- (2) The UTC and local time of the incident.
- (3) A unique value for the incident, the session identifier, and code. IBM service personnel use this code to identify how the incident was discovered.
- (4) If Telnet is used, the IP address and port of the secondary connection.
- (5) The row and column, in the 3270 display buffer, of the field where the overlay occurred. The offset is the offset in the 3270 display buffer.
- (6) The location in the outbound packet when the overlay occurred.
- (7) Up to 32 bytes of the outbound packet are displayed.

- (8)
The location in the inbound packet that caused the overlay.
- (9)
Up to 32 bytes of the inbound packet are displayed.
- (10)
The time stamp when the buffer was captured.
- (11)
The VTAM transmission and request headers.
- (12)
The RU data. The first and last sequence numbers of the RU chain that contributed to the RU are formatted.
- (13)
The time stamp when the trace date is recorded.

Using SMF

The 3270 IDS incidents are written to the System Management Facility (SMF) as a series of type 119 (subtype 81) records. Each record contains a common section that describes the incident and a saved DSCOUNT outbound buffer. The last outbound SMF record for an incident contains the inbound buffer.

For more information about the record, see [VTAM 3270 Intrusion Detection Services event record \(subtype 81\)](#) in *z/OS Communications Server: IP Programmer's Guide and Reference*.

Incident validation

When an incident is reported, it must be validated by gathering documentation immediately. This documentation should include the following information:

- The time and place that the incident occurred.
- The source, which is the logical unit (LU) names of the primary and secondary LUs. If the session is a TELNET session, the source also includes the IP address of the secondary LU.
- The name and type of application that was being used; and if possible, the transaction that was being executed.
- The name of the PU, LINE, and major node of the secondary LU, if applicable.
- Additional trace data needs to be collected to determine whether a pattern of data exists to this incident.

Example

```
IST2424I 3270 DATA STREAM ERROR - NETA.TS00002 NETA.TCPM0001
IST2425I PLU SUBAREA = X'0001' INDEX = X'0000' ELEMENT = X'0058'
IST2425I SLU SUBAREA = X'0001' INDEX = X'0001' ELEMENT = X'0009'
IST2441I JOBNAME = JHACKER SID = EAABEEC331E8DB02
IST2426I IPADDR = 192.168.98.254..61691
IST2427I DATE = 2016/01/25 TIME = 15:47:56 ID = 1
IST2428I ROW = 9 COLUMN = 16
IST2429I OUTBOUND - SEQ = X'0001' OFF = 598 LEN = 39
IST2431I 40404040 40404040 D1C1C3D2 E2D6D540 * JACKSON *
IST2430I INBOUND - SEQ = X'0001' OFF = 284 LEN = 39
IST2431I 40404040 40404040 F1F2F3F4 F5F6F7F8 * 12345678*
IST314I END
```

- The date and time of this incident is identified in message IST2427I and in the formatted trace data. The ID shows a unique identifier for this incident and this is the first one since VTAM was started.

```
IST2427I DATE = 2016/01/25 TIME = 15:47:56 ID = 1
```

- The secondary LU is identified in message IST2424I as NETA.TCPM0001. The following information displays this LU. Message IST271I shows that this LU is an application that the job named TELNET opens. Messages IST1727I and IST1669I identify the domain service name and IP address of the user.

Note: TCPM0001 is an application that acts as a secondary LU, which is not supported for 3270 IDS monitoring.

```
IST2424I 3270 DATA STREAM ERROR - NETA.TS00002 NETA.TCPM0001
d net,id=NETA.TCPM0001
IST097I DISPLAY ACCEPTED
IST075I NAME = NETA.TCPM0001, TYPE = DYNAMIC APPL 456
...
IST231I APPL MAJOR NODE = TCPAPPLS
IST271I JOBNAME = TELNET, STEPNAME = TELNET, DSPNAME = IST19405
...
IST1727I DNS NAME: JOEHACKER.FARFARAWAY.EXAMPLE.COM
IST1669I IPADDR..PORT 192.168.98.254..61691
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME          STATUS          SID          SEND RECV VR TP NETID
IST635I TS010002 ACTIV-P      EAABEEC331E8DB02 0004 0009      NETA
IST314I END
```

- The name of the PLU application is TSO0002. This user is logged onto TSO. The following information displays the application information. Message IST271I shows the TSO user ID. Messages IST2433I and IST2434I show the application 3270 IDS parameter values. Message IST2435I confirms that an 3270 IDS data steam error occurred.

```
IST2424I 3270 DATA STREAM ERROR - NETA.TS00002 NETA.TCPM0001
D NET,ID=TS00002,E
IST097I DISPLAY ACCEPTED
IST075I NAME = TS00002, TYPE = APPL 479
IST486I STATUS= ACT/S, DESIRED STATE= ACTIV
...
IST231I APPL MAJOR NODE = TS01A
IST213I ACBNAME FOR ID = TS010002
...
IST271I JOBNAME = JHACKER, STEPNAME = 0S390R5, DSPNAME = IST71E8A
...
IST2433I DSMONITR = YES, DSCOUNT = 15, DSACTION = (CONSOLE,NONE)
IST2434I DSTRUST = LOCALLU
IST2435I SESSIONS MONITORED = 1, ERRORS DETECTED = 1
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME          STATUS          SID          SEND RECV VR TP NETID
IST635I TCPM0001 ACTIV/E-S      EAABEEC331E8DB02 0009 0004      NETA
IST314I END

D NET,TSOUSER,ID=JHACKER
IST097I DISPLAY ACCEPTED
IST075I NAME = JHACKER, TYPE = TSO USERID 623
IST486I STATUS= ACTIV, DESIRED STATE= N/A
IST576I TSO TRACE = OFF
IST262I ACBNAME = TS00002, STATUS = ACT/S
IST262I LUNAME = TCPM0001, STATUS = ACT/S
IST1727I DNS NAME: JOEHACKER.FARFARAWAY.EXAMPLE.COM
IST1669I IPADDR..PORT 192.168.98.254..61691
IST2203I CHARACTER SET 02B9 - CODE PAGE 0417
IST314I END

D A,JHACKER
CNZ4105I 15.58.22 DISPLAY ACTIVITY 638
JOBS      M/S      TS USERS      SYSAS      INITS      ACTIVE/MAX VTAM      OAS
000000    00011    00002      00033      00003      00002/00300          00004
JHACKER  OWT      A=0025    PER=NO     SMC=000    PGN=N/A    DMN=N/A    AFF=NONE
              CT=000.032S  ET=01.04.21
              WUID=TSU00029
              WKL=TSO      SCL=TSO      P=1
              RGP=N/A      SRVR=NO     QSC=NO
              ADDR SPACE ASTE=1EFD6940
```

- The information of a secondary LU might identify the PU, LINE, and major node. In this example, the information of the PU, LINE, and major node is not available. However, you can use the TCPIP

commands **NSLOOKUP** and **TRACERTE** to confirm the ID and location of the secondary LU. Information about router206 indicates the approximate location.

For more information about TCPIP commands, see [z/OS Communications Server: IP System Administrator's Commands](#).

```
nslookup 192.168.98.254
EZB3170I Server: dns.example.com
EZB3172I Address: 192.168.100.4

EZB3170I Name: joehacker.farfaraway.example.com
EZB3172I Address: 192.168.98.254
READY
tracerte 192.168.98.254
CS V2R1: Traceroute to 192.168.98.254 (192.168.98.254)
1 router65.faraway.example.com (192.168.105.65) 2 ms 0 ms 0 ms
2 router1.faraway.example.com (10.6.0.1) 1 ms 0 ms 0 ms
3 router41a.faraway.example.com (192.168.120.41) 0 ms 0 ms 0 ms
4 routeredge201.faraway.example.com (192.168.106.201) 0 ms 0 ms
5 router1a.faraway.example.com (192.168.184.1) 15 ms 18 ms 21 ms
6 router8.faraway.example.com (192.168.34.8) 12 ms
7 router208.faraway.example.com (192.168.106.208) 2 ms 9 ms 11 ms
8 router12.faraway.example.com (192.168.96.120) 7 ms 12 ms 10 ms
9 joehacker.faraway.example.com (192.168.98.254) 2 ms 2 ms 1 ms
READY
```

- You can use the TCPIP **Netstat** command to show the time when the connection started.

For more information about TCPIP commands, see [z/OS Communications Server: IP System Administrator's Commands](#).

Tip: Information about the IP session is recorded in type 119 SMF records. Subtypes 1 and 2 contain information about the TCP connection. Subtypes 21 and 22 contain information about the TELNET connection. For TSO sessions, type 30 records contain information about the TSO user.

```
netstat all (port 55516
MVS TCP/IP NETSTAT CS V2R1      TCPIP Name: TCPCS      20:57:34
Client Name: TELNET             Client Id: 00000024
Local Socket: ::ffff:192.168.105.112..23
Foreign Socket: ::ffff:192.168.98.254..61691
BytesIn: 00000000000000002422
BytesOut: 00000000000000009580
SegmentsIn: 0000000000000000247
SegmentsOut: 0000000000000000320
StartDate: 01/25/2016           StartTime: 17:33:56
Last Touched: 20:47:56          State: Establish
...
Application Data: EZBTNSRV TCPM0001 TS010002 ET B
----
READY
```

- The following information of messages from IST2428I to IST2431I indicates the overlay in the 3270 data stream. Near row 9 and column 16 in the 3270 display buffer, a field that contains the string JACKSON is replaced by the string 12345678. Messages IST2429I and IST2430I show the respective PIUs where the fields can be found.

```
IST2428I ROW = 9 COLUMN = 16
IST2429I OUTBOUND - SEQ = X'0001' OFF = 598 LEN = 39
IST2431I 40404040 40404040 D1C1C3D2 E2D6D540 * JACKSON *
IST2430I INBOUND - SEQ = X'0001' OFF = 284 LEN = 39
IST2431I 40404040 40404040 F1F2F3F4 F5F6F7F8 * 12345678*
```

Tip: Message IST2431I shows part of the raw 3270 data stream, which might include different 3270 orders. The presence of the Start Field order (x'1D') might indicate that a field attribute has been overlaid, which might cause the incident report. Another order is the Start Field Extended (x'29'). For more information about the 3270 data stream, see [3270 Data Stream Programmer's Reference](#).

- The following generalized trace facility (GTF) trace data shows information about the buffers. Start additional traces of VTAM buffers to verify whether the sequence is repeated. The TCPIP packet trace data can also be collected. The TELNET option of the TCPIP packet trace formatter can be used to display the 3270 data stream orders.

For more information about the TCPIP packet trace, see [z/OS Communications Server: IP Diagnosis Guide](#).

```

(11)VTAM      TH=40000000 00000000 00010001 00000001 1800000B 00580001 051F RH=0380C0
(12)      SEQ 0001-0001      F5C21140 402901C0 40F4F040 40E44040 40404040 *5B.  ..{ 40
U      *
      404040C3 C8D9C9E2 E3C9C1D5 40404040 40404008 *
CHRISTIAN      .*
...
HS M *      114DC829 01C0E9C5 F94040D7 40C8E240 40D44040 *. (H..{ZE9 P
      40D4C1E2 D6D54040 40404040 40404011 4DF02901 *
MASON      . (0..*
      C06CF6C3 4040D740 4040C940 40404040 D1C1C3D2 *{ %6C P
I      JACK*
      E2D6D540 40404040 40404040 114ED829 01C06DF6
*SON      .+Q..{_6*
...
      40404040 40404040 40C8C5E7 E2E3D9C9 D5C74DF0 *
HEXSTRING(0*
      F05D4011 5D7E1D60
*0) .)=. -      *
      (11)VTAM      TH=40000000 00000000 00000001 00010001 1C000058 000B0001 0298 RH=0393A0
      (12)      SEQ 0001-0001      7D4AD811 40E9C3F1 4040E440 40404040 D4404040 *'¢Q. ZC1
U      M *
      C1D3C5E7 E8E24040 40404040 40404040 11C1F9C3
*ALEXYS      .A9C*
      F54040E4 4040E240 40D44040 40D4C1E2 D6D54040 *5 U S M
MASON *
      40404040 40404040 4011C3C9 C3F94040 E440C8E2
*      .CIC9 U HS*
...
      40E4D540 40C940D4 404040D4 C1E2D6D5 40404040 * UN I M
MASON *
      40404040 40404011 4AC1F6F0 4040D740 40404040 *      .¢A60
P *
      40404040 F1F2F3F4 F5F6F7F8 F9404040 40404040 *
123456789      *
      114AE9F6 F14040D7 40404040 40D44040 40D4C1C4 *. ¢Z61 P
M MAD*
      C9E2D6D5 40404040 40404040 40114BF9 C5F54040
*ISON      . .9E5 *
...
      C8C540E5 C1D3E4C5 40E3D67A 40404040 40404040 *HE VALUE
TO:      *
      40404040 404040C8 C5E7E2E3 D9C9D5C7 4DF0F05D *
HEXSTRING(00)*
      40
*      *
```

Logon mode parameters

If you have an application program that acts as the secondary end of the session, you might need to create new entries or supplementary tables, as follows:

- If the application program is the secondary logical unit, you might want to specify session protocols. You can direct VTAM to the proper session parameters by specifying the logon mode table entry (through the DLOGMOD operand) and the logon mode table (through the MODETAB operand) on the APPL definition statement of the program.
- If the application program is intended to be in session with specific devices or application programs and it requires certain protocols, you might have to create new logon mode table entries associated with the devices.

Both an APPN Class of Service and a subarea Class of Service, specified in the logon mode entry, can be associated with each session. The APPN Class of Service determines the route that is used for APPN links, and the subarea Class of Service determines the virtual route that is used for a session. Therefore, you should understand the Class of Service requirements for your application programs. Considerations such as the importance of the function provided by the application program, the nature of the data exchanged with the logical units, and the expected response times contribute to determining the Class of Service needed for sessions with the application program. The Class of Service names must then be specified in

the logon mode table entries used for LU-LU sessions. For more information about logon mode entries, see [“Selecting session parameters for the logon mode table”](#) on page 203.

Using user variables (USERVAR)

The VTAM user variable (USERVAR) maps a generic application program name used in a terminal logon to the name of a specific application program, based on the value of the USERVAR. This function is used by IMS and CICS XRF to map user logons to the IMS or CICS application program that is currently active.

You can also use USERVARs in other ways, such as to facilitate migration from one application program release to another. The VTAM user variable:

- Enables your installation to use the USERVAR function without having to rely upon NetView command lists or manual procedures to propagate USERVAR updates. VTAM communicates USERVAR values across domains in the same or different networks.
- Enables your application programs to use USERVAR names across VTAM application programming interface (API) in place of logical unit names without requiring code changes.
- Enables your terminals to take advantage of USERVAR translation without your having to code interpret tables for them.

The VTAM application programming interface (API) allows application programs to use USERVAR names across the API in place of LU names. Because USERVAR can be used across the API, many application programs that participate in sessions with an XRF-capable application program (such as IMS or CICS) can remain unchanged when USERVARs are used.

Also, because VTAM determines whether a name refers to an LU or to a USERVAR and performs the appropriate translation automatically, you do not have to code an interpret table to specify that a name used in a particular logon is actually a USERVAR.

Note: The generic name used by USERVAR should not be used by ALIAS. An alias CDRSC should not exist in a host where USERVAR mapping is used unless the host is in the same network as the real resources.

Application workload balancing with USERVAR

A group of applications can be given a common or generic name. VTAM uses this generic name, known as a USERVAR, to associate a logon request with the currently active member of this group. This capability can be used for workload balancing and preparing for the planned takedown of a host. USERVAR values can be managed by the user and by VTAM.

For example, assume that the USERVAR, IMS, has been defined, and that there are two IMS application programs: IMS1 and IMS2. The generic name (USERVAR) can be assigned the value of either IMS1 or IMS2, based on which system is active. If the USERVAR for the IMS application programs is set to IMS1, for instance, then any session requested with IMS is routed to IMS1. If IMS is set to IMS2, then new sessions are routed to IMS2. Therefore, if the workload going to IMS1 is very heavy, the IMS USERVAR can be set to IMS2 and subsequent sessions are routed to IMS2.

The following commands display and modify the VTAM application names associated with USERVARs:

- DISPLAY USERVAR
- MODIFY USERVAR
- DISPLAY ID
- DISPLAY SESSIONS

For further information about these commands, see [z/OS Communications Server: SNA Operation](#).

Classes of USERVARs

VTAM supports the following two classes of USERVARs:

- User-managed USERVAR
- Automatic (or VTAM-managed) USERVAR

User-managed USERVAR

This class is explicitly set using the MODIFY USERVAR operator command. The MODIFY USERVAR command can be issued either manually by a human operator or automatically by a NetView command list or a VTAM application program (as is done by IMS or CICS for XRF).

VTAM does not attempt to change the value of a USERVAR that has been set explicitly using the MODIFY USERVAR operator command, so existing command lists and procedures for propagating USERVAR changes across domains in a network should continue to work as is. However, a MODIFY USERVAR command can be used to delete a user-managed USERVAR to subsequently allow that USERVAR name to be managed by VTAM.

At least one host in the network must have a user-managed USERVAR from which other hosts can obtain the USERVAR value and type. In XRF situations, IMS and CICS both set the USERVAR value at the site of the XRF active application program.

Notes:

1. For XRF application programs, VTAM deletes user-managed USERVARs when the application program terminates or enters a specific state, such as CONCT or INACT.
2. In a subarea-only network, the user-managed USERVAR can be defined on the VTAM that owns the real destination resource, on the VTAM where the session originates, or on any VTAM along the session path. However, because there might be attempts from many different VTAMs to establish sessions with the destination resource using the USERVAR name, the best results are achieved when the user-managed USERVAR is defined on the VTAM that owns the real resource, or on a VTAM that is as close to the owning VTAM as possible.
3. In an APPN network, or in a mixed APPN and subarea network, the location of the user-managed USERVAR is more restrictive than in other networks. If the real resource resides on a VTAM end node (EN), a network node (NN), or a migration data host (MDH), the user-managed USERVAR must be defined on that VTAM (the owning VTAM). Therefore, when the real resource is on an EN, define the USERVAR on the EN but do not define it on its network node server (NNS). If the real resource resides on an interchange node (ICN), or on a VTAM located in or through a subarea network attached to an ICN, the user-managed USERVAR must be defined on the ICN or on one of the VTAMs along the subarea portion of the session path.

When the MODIFY USERVAR command is issued on an EN, the USERVAR name is registered to the NNS of the EN unless the value UVEXIT=NO is coded (or is the default value) and the real resource name is either a dynamic CDRSC or is not defined.

Automatic (or VTAM-managed) USERVAR

This class is created automatically by VTAM in one domain as a copy of a user-managed USERVAR of the same name in another domain. You do not have to do anything to allow VTAM to create automatic USERVARs. When a VTAM SSCP sends a cross-domain session-initiation request to another domain and receives a reply indicating that the name specified in the original session request is really a USERVAR with a particular value and not the name of a logical unit, the SSCP creates an automatic (VTAM-managed) USERVAR to save the information. The cross-domain session request is then resent using the real LU name, which was passed back in the original reply.

Types of USERVARs

Changes to the value of a USERVAR are not propagated by VTAM at the time of the change. Instead, when a domain using a USERVAR to establish a session encounters a specific USERVAR for the first time, it uses an adjacent SSCP table to perform a cross-domain search to obtain the USERVAR value. Subsequent attempts to use the same USERVAR can result in a repeat of the cross-domain search to update the USERVAR value, depending on the value of the TYPE operand. The TYPE operand is specified on the MODIFY USERVAR command when the USERVAR is created. The TYPE operand is copied, along with the USERVAR value, when an automatic USERVAR is created in another domain. Following are the types of USERVARs:

Static USERVAR

The value is assumed not to change. It is saved as an automatic USERVAR upon initial use but is not updated automatically using the cross-domain search process (unless the existing copy of the automatic USERVAR is explicitly deleted by the operator).

Dynamic USERVAR

The value is expected to remain relatively stable but change occasionally (for example, when an XRF takeover occurs). To establish a cross-domain session to an application program identified by a dynamic USERVAR, VTAM must recheck the USERVAR value after every abnormal session termination involving the application program that it references.

DYNAMIC is the default value for the MODIFY USERVAR command and is recommended for use within XRF complexes.

Volatile USERVAR

The value is expected to change often. Rather than creating an automatic copy of the USERVAR, a remote VTAM repeats the cross-domain search to reestablish the USERVAR current value each time it is referenced to establish a session. Because of the additional search flow that is necessary for each session establishment, use of volatile USERVARs might have a substantial impact on the time required for session establishments.

Note: If a USERVAR value is allowed to vary over multiple active application programs, LU 6.2 communication works only for a recovery environment. See [“LU 6.2 in an XRF Environment”](#) on page 326.

Processing USERVARs

Assume that an XRF application program, such as CICS, runs on Systems A and B of the sample network shown in Figure 94 on page 321. CICS_A is the specific VTAM APPLID of the CICS running on System A, and CICS_B is the APPLID of the CICS running on System B. The terminal network is owned by the VTAM in System C (this is a communication management configuration, or CMC).

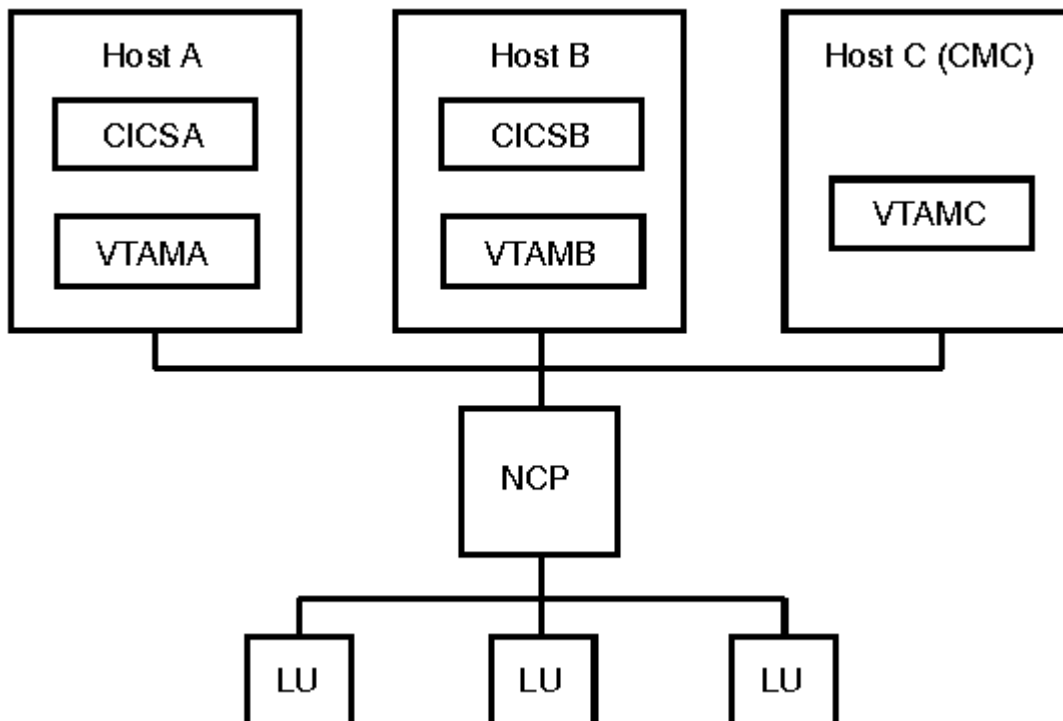


Figure 94. Sample extended recovery facility network

1. When CICS is started as an active CICS application program, it issues an `F NET, USERVAR, ID=CICS, VALUE=CICSA` command, requesting that VTAMA create a dynamic (by default), user-managed USERVAR with a value of CICSA. Because CICS is started as an XRF alternate, it does not request that VTAMB create a USERVAR.
2. Logon requests from a terminal are processed by VTAMC, which is the controlling VTAM for the terminal network. The first time a terminal attempts to log on to generic application name CICS, VTAMC processes the logon request as follows:
 - Because VTAMC does not recognize the name CICS yet, it treats the request like any other cross-domain logon request and sends a cross-domain session request to other VTAMs listed in an ADJSSCP table looking for CICS. At this point, VTAMC does not know whether CICS is the name of a USERVAR or a real VTAM APPLID.
 - When VTAMA receives the session request, CICS is identified as the name of one of its own user-managed USERVARs. It replies to VTAMC, indicating that CICS is really a dynamic USERVAR whose value is currently CICSA.
 - VTAMC uses the information it received from VTAMA to establish a cross-domain session with the application program by sending another session request using that specific application name, CICSA. In addition, VTAMC creates an automatic USERVAR to save the USERVAR information for future use.
3. For subsequent terminal logons to CICS, VTAMC uses the information it saved previously in the automatic USERVAR to translate the requests immediately to CICS. VTAMC continues to translate CICS requests directly to CICS without repeating the cross-domain search process because CICS is a dynamic USERVAR. This process continues until either an XRF takeover occurs or VTAMC discovers that it is otherwise unable to establish sessions with CICS.
4. If CICS fails and CICSB takes over and becomes the new active application program, CICSB issues an `F NET, USERVAR, ID=CICS, VALUE=CICSB` command, requesting that VTAMB create a new user-managed USERVAR with the value CICSB. This means that VTAMB is now the VTAM to identify session requests searching for CICS. VTAMB then responds, indicating CICSB is the current value of the USERVAR.

Each VTAM that was notified that its XRF sessions were switched deletes any automatic USERVARs that were used to create those sessions. This means that the next logon for CICS does not find the USERVAR and goes through the cross-domain search process (step “2” on page 322) to determine the new value of CICS. In addition, the VTAM associated with the failing application program (VTAMA) deletes its user-managed USERVAR for CICS and prevents any further logons to CICS to ensure that logons do not get directed to the failing application program.

Note: If a resource submits a network-qualified session request with a generic resource name, USERVAR processing occurs in the domain of the resource if one of the following is true:

- The NETID specified on the session request is equal to the NETID of the USERVAR value.
- The NETID specified on the session request is equal to the NETID of the host and the USERVAR value is not network-qualified.

USERVAR propagation and routing

NetView command lists are not required to propagate USERVAR modifications to a host in which VTAM can manage the USERVAR values; however, existing NetView USERVAR propagation command lists continue to work as they did before.

Note: VTAM can manage USERVAR values at a given host only if it and all hosts on the session establishment path have VTAM Version 3 Release 2 or later and have installed the appropriate PTF for VTAM Version 3 Release 2 that provides this USERVAR support.

VTAM searches for the value of a USERVAR through ADJSSCP table routing for a cross-domain session request. Therefore, if you have not already done so, you might need to define adjacent SSCP tables for searches to find the VTAM having the user-managed USERVAR. For information about this, see [“Adjacent SSCPs” on page 431](#).

Generally, VTAM does not change user-managed USERVARs; they are viewed as being the responsibility of the user (or application program) that created them. However, in XRF situations, VTAM reduces the problem of session requests going to a failing XRF-active application program by deleting the user-managed USERVAR that references the failing application. VTAM deletes the user-managed USERVAR only at the site of the failing application program. If you have created a user-managed USERVAR at another host in the network, you need to make sure that USERVAR is deleted or that it continues to direct session requests to the failing application program.

For the VTAM having a user-managed USERVAR to respond to a session initiation request using that USERVAR, VTAM must know certain information about the LU that its value references (for example, the name of the SSCP that owns the LU). The necessary information is available if the referenced LU is owned by the host having the user-managed USERVAR, or if a CDRSC exists for the LU on that host. Otherwise, VTAM responds to the session initiation request as though the name were not found to allow the ADJSSCP routing to continue.

Defining your network with USERVARs

When you define your network resources for XRF, you should consider the following situations:

- If you allow dynamic CDRSCs and trial and error rerouting, you do not need to predefine a CDRSC for the USERVAR name or CDRSCs for the active and backup application program.

Note: Referencing dynamic CDRSCs with USERVARs can cause situations in which session initiations fail intermittently, because dynamically defined CDRSCs can be deactivated and deleted by VTAM, depending on the dynamic CDRSC timer value (CDRSCTI) and whether there are any active sessions.

- If you do not allow dynamic CDRSCs to be created in the network, you need to predefine a CDRSC for the USERVAR name, a CDRSC for the active application program, and a CDRSC for the backup application program. If you do predefine a CDRSC with the USERVAR name, you should not specify the cross-domain resource manager (CDRM) name because the USERVAR may not always be managed from the same SSCP. However, you should define the CDRSC under the network in which the USERVAR resides to reduce the amount of time for the ADJSSCP search process.

The following applies regardless of whether you use USERVARs with XRF:

- You can use USS tables, interpret tables, or both with USERVAR translation; the USS and interpret processing is done first and USERVAR translation is done using the output of those processes.
- You are not required to define interpret table entries for terminals to use the USERVAR function. If you have already done so, however, you do not have to remove them.
- If the application is cross-domain, the USERVAR name must be defined as a CDRSC using either static or dynamic definitions and must be unique within the network.
- The OLU host of a session initiation and the host that is resolving the USERVAR name must support name translation. For VTAM 4.1 and earlier, the TRANSLAT start option must include USERVAR. In VTAM 3.4.2 and later, RACALIAS must include USERVAR in ISTRACON.

Dynamic USERVAR update session failure

Whenever an abnormal session termination occurs for a session that was initiated using a USERVAR, VTAM deletes the USERVAR if it is VTAM-managed and dynamic. When the next session request arrives at any host other than the XRF primary host, VTAM repeats the search for the USERVAR and saves the new current value in an automatic USERVAR.

The following kinds of session termination are considered normal because they do not cause the resetting of a dynamic USERVAR:

- Requests for conditional or unconditional termination, such as:
 - Operator issued the VARY NET,TERM,TYPE=COND|UNCOND command.
 - SLU application program issued the TERMSESS OPTCD=COND|UNCOND command.
 - Terminal issued (using USS) the LOGOFF TYPE=COND|UNCOND command.

- Session setup failures because of the following sense codes:
 - 08210002
 - 088A0003
 - 088A0004
- LU-LU session terminations (UNBIND) for the following reasons:
 - Normal end of session (TYPE=X'01')
 - CLSDST PASS (TYPE=X'02')
 - Session parameters that are not valid (TYPE=X'06')
 - Other UNBINDs between X'01' and X'06'
 - Format or protocol error (TYPE=X'FE')

A session termination for any reason other than those reasons described previously is handled as an abnormal session termination and causes the deletion of any VTAM-managed dynamic USERVAR associated with that session. Following are examples of abnormal terminations:

- Requests for forced termination which, like requests for conditional or unconditional termination, can come from operators, SLU application programs, or terminals.
- LU-LU session terminations, including:
 - VR INOP (TYPE=X'07')
 - Route extension INOP (TYPE=X'08')
 - LU failure (TYPE=X'0C')
 - Cleanup termination, such as V NET,INACT (TYPE=X'0F')
 - Gateway node cleanup (TYPE=X'11')
- Session setup failures other than those previously described (such as failure to find the destination LU after using a dynamic USERVAR to determine the application program name). In such cases, VTAM deletes the dynamic USERVAR and tries the session initiation again using the USERVAR name to determine the current value of the user-managed USERVAR.

Whenever the active application program in an XRF complex fails, notification of the session failures is sent to the SSCPs of the session partners, thereby causing them to reset their dynamic USERVARs. In the host of the XRF active application program, VTAM deletes all USERVARs that reference the failing application program to prevent future USERVAR searches from using a USERVAR that references the failing application program.

If a backup session already exists, an alternate application program sends a switch request to the NCP to takeover the sessions. The XRF takeover can be initiated before VTAM even knows that the active application program is failing. The VTAM first awareness of such a takeover is the receipt of a session ended notification having a specific session termination reason code (X'13'), which indicates that the primary session terminated because of an XRF switch. At a remote SSCP, such session terminations are handled as abnormal session terminations. At the SSCP of the failing active application program, however, the SSCP also does the following actions:

1. Marks the application program as "not enabled for sessions" to prevent any further session establishments
2. Deletes the USERVAR that is associated with the session that is ending

Generic resources function

Whereas a USERVAR is assigned to only one resource at a time, the generic resources function allows the assignment of a generic resource name to a group of active application programs that all provide the same function. The generic resource name is assigned to multiple active application programs simultaneously, and VTAM automatically distributes sessions among these application programs rather than assigning all sessions to a single resource. Thus, session workloads in the network are balanced. Session distribution

is transparent to users; an LU initiates a logon request using the generic resource name and need not be aware of which particular application program is actually providing the function.

The generic resources function also increases application program availability, as each active application program that uses a given generic resource name (a generic resource member) can back up other generic resource members. Thus, no single application program is critical to resource availability. When a generic resource member fails, an LU can reinitiate its session using the same generic resource name. VTAM resolves the session initiation to one of the other generic resource members. Because the user is unaware of which generic resource member is providing the function, the user is less affected by the failure of any single generic resource member.

In the same way, an application program can be added as a generic resource member and accessed using the same generic resource name. Because the user is unaware of which generic resource member is providing the function, an application program can be added as a generic resource member with little or no impact to the application user. At the same time, the additional application program provides an immediate improvement in performance and availability because the workload is now shared by an additional application program.

Use of generic resources requires a coupling facility structure. See [“Generic resources” on page 352](#) for additional information concerning generic resource requirements.

High availability using extended recovery facility

You can identify an application program as capable of operating in an extended recovery facility (XRF) environment. Two IBM application programs, Information Management System (IMS) and Customer Information Control System (CICS), support an XRF environment.

XRF provides an alternate application program when certain host components fail. It reduces the impact of planned outages or outages caused by failures in VTAM, MVS, IMS, CICS, or the host processor on selected users of IMS, CICS, or other application programs that use XRF functions. The types of failures that cause the takeover of an XRF session are determined by how IMS, CICS, or the application program using XRF is defined. XRF might not reduce the impact of failures in the NCP or in telecommunication lines, depending on which NCP or line fails. However, the XRF function in CICS or IMS does not automatically perform a switch to the alternate application program under these conditions. Therefore, either the operator must switch from the active to the alternate system, or you can use NetView automation to perform the switch.

To enable high availability, use the HAVAIL operand on the APPL definition statement.

VTAM, in conjunction with the application program, establishes a primary session from the currently active application program and a backup session from the alternate application program to the terminal. The NCP for a terminal in an XRF environment maintains two sessions through two VTAMs: one to a primary XRF application program and one backup session to the alternate XRF application program. The XRF application programs, in conjunction with VTAM and NCP, switch to the backup session when some component fails. The active and alternate application programs must be in the same network. However, a terminal participating in an XRF session can be in a different network than either of the application programs.

One application program acts as the active application program, and the other acts as the alternate in case the active application program or another component in the active host fails. Both application programs are identified by a single name that is associated with the currently active application program.

Although one VTAM in a single-processor environment can support both a primary and a backup session, this configuration does not provide the availability characteristics usually associated with an XRF environment.

XRF is also supported for application programs that use LU 6.2 sessions to communicate with other applications or logical units in peripheral nodes attached to an NCP. For further information, see [“LU 6.2 in an XRF Environment” on page 326](#).

Security features in an XRF environment

An XRF application program can establish cryptographic sessions and message authentication sessions with other LUs. Cryptography protects data passing over lines by permitting enciphering and deciphering of data for LU-LU sessions. Message authentication provides a message authentication code used to validate the contents of the data.

The XRF application program must reside in a host that has the IBM Integrated Cryptographic Service Facility/MVS (ICSF/MVS) or a compatible cryptographic product installed and active.

The following references are used with compatible cryptographic products::

PCF/CUSP

Refers to any cryptographic product that is compatible with PCF/CUSP.

CCA

Refers to any product that is compatible with Common Cryptographic Architecture (CCA).

Many PCF/CUSP compatible cryptographic products must be started before starting VTAM. They also must be started before you activate an external CDRM or CP for which CROSS statements have been defined in the cryptographic key data sets. This is not necessary for the VTAM Integrated Cryptographic Service Facility.

The XRF application program can specify that cryptography is either selective (specified by the session end) or required (session ends must support cryptography).

To enable this facility, code the following statements:

- ENCR operand on the APPL definition statement for the XRF application program.
- ENCR operand on the LU definition statement for any logical units that will be session ends.
- ENCRTYPE keyword can be coded on both the APPL and LU. To enable TDES24 encryption, ENCRTYPE=TDES24 must be coded; otherwise the ENCRTYPE will default to DES.

Note: A corequisite of NCP Version 7 Release 8 is required for XCF/Crypto with triple-DES operation.

- Cryptographic keys in the cryptographic key data set.

For more information about coding the ENCR operand, see [“Cryptography facility” on page 289](#). For more information about coding cryptographic keys, see [Appendix E, “Cryptographic keys,” on page 593](#).

LU 6.2 in an XRF Environment

USERVAR process of using a generic name to represent a service does not work in an active LU 6.2 environment because LU 6.2 cannot tolerate LU name substitutions. However, support for LU 6.2 in an XRF environment allows generic name substitution, but only in a recovery environment.

The LU 6.2 application uses a USERVAR ID during recovery to transfer its users from one application program to a backup application program. This provides faster backup for failing applications.

To use this support, you need to identify the USERVAR to be used for LU 6.2 communications.

Be aware of the following restrictions:

- You can assign only one USERVAR ID to an LU 6.2 application program.
- You can have only one LU 6.2 application program as the primary application program to provide a service.
- The operator can modify the USERVAR ID for LU 6.2, but the command must be issued from the host system of the application program providing the service.
- If an incorrect USERVAR ID is assigned to an active LU 6.2 application program, the ACB must be closed and then reopened with the correct USERVAR ID.
- Even sessions that request a bind using the real LU name must still carry the USERVAR ID.

Persistent LU-LU sessions

Application programs supporting persistent LU-LU sessions can reestablish LU-LU sessions at the level before the failure without the session establishment flow. There are two levels of support for persistent LU-LU sessions.

- Single node persistent sessions (SNPS)

If an application program fails or is brought down, VTAM can retain active sessions, allowing the same or another application program to reconnect to the sessions, avoiding the need to reestablish the sessions. If an application program fails, it can reconnect to the retained sessions when it recovers. Also, another application program can take over the sessions. SNPS takeover processing allows an application program to take over the sessions of an active application, although the active application is able to indicate, through SETLOGON PERSIST, that such requests are rejected by VTAM.

- Multinode persistent sessions (MNPS)

Multinode persistent LU-LU sessions (MNPS) extends the single node persistent session support to application, VTAM, operating system, and hardware failures. Applications can be restarted on an alternate VTAM and their sessions restored after a VTAM failure. Applications can also restart on the same VTAM after VTAM has been restarted.

MNPS planned takeover processing differs from SNPS takeover processing in that the application is required to fail, or issue CLOSE ACB, while enabled for persistence in order to be eligible to move to another VTAM.

An alternative to MNPS planned takeover, called MNPS forced takeover, works closer to SNPS takeover processing, in that the target application does not have to fail or issue a persistent CLOSE ACB before the MNPS takeover attempt being made. However, unlike SNPS takeover, MNPS forced takeover requires that the taking over application indicate on the ACB macroinstruction that any OPEN ACB attempted using this ACB is an MNPS forced takeover attempt; likewise, the target application must indicate on SETLOGON OPTCD=PERSIST that it will accept MNPS forced takeovers. SNPS takeover processing does not require any special communication from the application, but the capability does exist for the application to indicate that SNPS takeover is not permitted.

When an application program is persistent enabled, VTAM saves data sent and received by the application program and other status information. VTAM uses this information to reestablish the session after a failure has occurred, or as part of planned or forced takeover processing.

For an application to be capable of persistence, code PERSIST=YES on the ACB macroinstruction. To enable persistence, the application program must specify OPTCD=PERSIST on the SETLOGON macroinstruction. The application program can also disable persistence by issuing SETLOGON OPTCD=NPERSIST at any time. An application that wants to use multinode persistent session support must have PERSIST=MULTI coded on its application definition statement.

For an OPEN ACB to be considered for MNPS forced takeover processing, include FORCETKO=YES on the ACB macroinstruction (along with PERSIST=YES).


The application is permitted to specify the following types of forced takeover requests it will accept from other application images of the same name:

- To indicate that the application will accept MNPS forced takeover requests from other instances of the application, code PARMS=(FORCETKO=ALL) or PARMS=(FORCETKO=MULTI) on the SETLOGON OPTCD=PERSIST macroinstruction.
- To indicate that the application will accept SNPS forced takeover requests from other instances of the application, code PARMS=(FORCETKO=ALL) or PARMS=(FORCETKO=SINGLE) on the SETLOGON OPTCD=PERSIST macroinstructions.
- To prevent any forced takeover from being processed for the application, code PARMS=(FORCETKO=NONE) on the SETLOGON OPTCD=PERSIST macroinstruction.

Rule: Issuing SETLOGON OPTCD=NPERSIST does not affect the setting of the FORCETKO capability of the application.

Single node persistent sessions


When a persistence-enabled application program fails, VTAM retains the sessions, saves the allocated resources and control blocks, and shields the network from knowledge of the application program failure. VTAM stores the incoming data so that the network views the session as active but not currently responding. When the failed application program restarts or another application program takes over, VTAM reconnects the sessions. The backup of a persistence-enabled application program is shown in the following figures:



1. APPL1 has initiated a session with terminal T. The session is persistent capable and enabled. APPL1 specified persistence on both the OPEN ACB (PERSIST = YES) and the SETLOGON (SETLOGON OPTCD = PERSIST) macroinstructions.

APPL1 has initiated a session with terminal T. The session is persistent capable and enabled. APPL1 specified persistence on both the OPEN ACB (PERSIST = YES) and the SETLOGON (SETLOGON OPTCD = PERSIST) macroinstructions.


Figure 95. Application program backup using persistent LU-LU sessions - part 1



2. APPL1 fails. The session is pending recovery. If no timer is specified (PSTIMER), the session can remain pending recovery indefinitely. However, the retained session can be terminated by a VARYNET INACT command.

APPL1 fails. The session is pending recovery. If no timer is specified (PSTIMER), the session can remain pending recovery indefinitely. However, the retained session can be terminated by a VARY NET INACT command.

Figure 96. Application program backup using persistent LU-LU sessions - part 2



3. The operator or program operator detects that APPL1 has failed and uses APPL1' as a backup. APPL1' issues an OPEN macroinstruction with the same ACB name APPL1 issued. It then restores the session. Takeover is transparent to terminal T.

The operator or program operator detects that APPL1 has failed and uses APPL1' as a backup. APPL1' issues an OPEN macroinstruction with the same ACB name APPL1 issued. It then restores the session. Takeover is transparent to terminal T.

Figure 97. Application program backup using persistent LU-LU sessions - part 3

For more information about enabling an application program to support persistent LU-LU sessions, see the [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#) and [z/OS Communications Server: SNA Programming](#).

The HALT and VARY INACT,TYPE=FORCE commands override single node persistent sessions.

VTAM common network services

When you use VTAM common network services (VCNS) with the ES/9000 series communication adapter to send data across an X.25 packet switched data network, specify the VCNS=YES operand on the APPL definition statement. This allows sharing the physical connectivity of the communication adapter between SNA and VCNS data traffic.

You can use VCNS with the IBM 3172 Nways interconnect controller to send data to a local area network. Specify VCNS=YES on the APPL definition statement to allow sharing the physical connectivity of the communication adapter between SNA and VCNS data traffic.

Cross-memory application programming interface (API)

Sessions can be established by an application program running in an address space other than that which opened the application ACB of the program. This enables the application program to be represented to VTAM by one ACB (that is, APPL definition statement) but span multiple address spaces. The application program can also issue VTAM requests to send and receive data while running in a different address space than the one in which the session was originally established.

For information about enabling an application program to support cross-memory API, see [z/OS Communications Server: SNA Programming](#).

Allocating private storage

All incoming data is queued in a VTAM-owned ESA data space, not in private storage. The MAXPVT operand therefore has no effect on data queued for an application.

Communicating with start-stop devices

When communicating with a start-stop device that is supported by the Network Terminal Option (NTO) or the X.25 NCP packet switching interface (NPSI), special data stream characteristics of the non-SNA terminal might need to be specified for use by an application program. You indicate these characteristics

by specifying the appropriate value on the TERM operand of the PU or LU definition statement that defines the virtual physical or logical unit. The application program can get information about these characteristics by using the INQUIRE OPTCD=DEVCHAR macroinstruction. It can also get it from the CINIT RU, which is queued when a logon is received from one of these devices.

LU 6.2 application programs

LU 6.2 enables a VTAM application program to communicate with nonhost application programs on an advanced program-to-program communication (APPC) level. Without LU 6.2, the nonhost program usually must emulate a terminal when communicating with a VTAM application program, using protocols that are not well-suited for application programs and that diminish the nonhost application program ability to handle errors and to initiate and end communications.

LU 6.2 protocols are designed to optimize use of sessions and improve the efficiency of data transfer through the network. LU 6.2 can also simplify the program logic needed to establish communications with other application programs and to send and receive data. Though VTAM support simplifies the task of writing application programs that use LU 6.2, be aware that enabling VTAM support requires some LU 6.2 functions to be implemented by an application program.

You can write or modify host application programs to use LU 6.2 protocols throughout the network.

Enabling LU 6.2 support

You enable LU 6.2 support for an application program by specifying the APPC=YES operand on the APPL definition statement. The application program must use the APPCCMD macroinstruction to send and receive data on LU 6.2 sessions and must allow VTAM to manage LU 6.2 session establishment and termination. This means that an application program cannot use non-LU 6.2 VTAM macroinstructions such as OPNDST, SEND, or RECEIVE on LU 6.2 sessions; however, the same application program can use non-LU 6.2 macroinstructions to communicate on non-LU 6.2 sessions. For information about macroinstructions for LU 6.2 sessions, see the [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

Existing application programs that have implemented LU 6.2 without VTAM APPCCMD support (and consequently use the restricted non-LU 6.2 macroinstructions) should not have APPC=YES on their APPL definition statement. They do not work if it is included. To take advantage of VTAM support, they must be migrated to use the APPCCMD macroinstruction.

LU 6.2 sessions

Following is information about how LU 6.2 sessions operate.

Session establishment and termination

VTAM handles the tasks of session establishment and termination for LU 6.2 sessions for those application programs that use LU 6.2 support. Application programs explicitly issue conversation establishment and termination requests. VTAM automatically sets up and takes down sessions in response to these requests, according to the session limits the application programs have requested and negotiated. The application program's main involvement with session management is in establishing limits on the number of sessions that it can have with another LU in the network.

Conversations

Sessions between LU 6.2 session partners are used by transaction programs that operate in each partner LU to perform data exchange. Rather than sessions, the basic communication link between logical units using LU 6.2 protocols is a conversation. A conversation is the exclusive use of a session between two partner LUs for some unit of work. Sessions between LU 6.2 session partners are serially reusable. When one conversation ends, another conversation can use that same session between the partner LUs.

Sending and receiving data

If a session is full-duplex, conversations on that session can be full-duplex or half-duplex. If a session is half-duplex, conversations on that session must be half-duplex. In full-duplex conversations, an LU can send and receive data simultaneously. In half-duplex conversations, only one LU can send data at a given time. One transaction program in the conversation must be designated as the sender and the other as the receiver. Sessions between logical unit partners are designated as contention-winner or contention-loser sessions. If the two LU partners attempt to use a session simultaneously, the designated contention winner for the session has priority for the use of the session. An application program can allocate a conversation over a session for which it is designated the contention loser, but the logical unit partner must give the application program permission.

Session characteristics

Session characteristics are determined by the mode name associated with the session. Each session between LU 6.2 partners is related to a specific mode name. The mode name represents a set of characteristics associated with that session. All other sessions between the LU 6.2 partners that use that mode name are assumed to have the same set of session characteristics.

Session initiation using resource verification reduction

In an APPN network, LU-LU session initiation requires that the control point serving the requesting LU obtain partner LU location and route information to establish the session. For example, when the control point in the end node of the requesting LU receives the session initiation request, it forwards a Locate request to its network node server to locate the partner LU. If the resource is known to the network node server, the network node server sends a verification search directly to the resource. If the resource is not known to the network node server (that is, if the resource information has not been previously cached), the network node server conducts a network-wide search for the resource.

After information about a resource has been cached, the resource verification reduction function can be used to eliminate the verification search for the partner LU. Either the network node server of the requesting LU or the network node server of the partner LU may use the cached information to build a reply on behalf of the destination LU when resource verification reduction is requested. As long as the cache information is accurate and complete, session setup time is reduced. The maximum amount of time cached information will be used for resource verification reduction is controlled by the VFYREDTI start option. If the information is older than the value of the VFYREDTI start option value, resource verification will be performed upon receiving the next request for the resource. For more information about the VFYREDTI start option, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Cached information that is not accurate may cause the session setup to fail. When this occurs, VTAM will redrive the session initiation request, forcing verification of the resource.

If the requesting application has a LOGON exit, it is driven for a second time when the second initiation request is successful. However, the CINIT information for the second request contains different session information than that of the initial LOGON exit invocation. Thus, an application program with a LOGON exit must be able to have its LOGON exit driven multiple times as a result of one logical request (for example, an ALLOCATE). Application programs that support having their LOGON exits driven multiple times must indicate this support, by coding the application capabilities vector with the multiple logon exit support indicator on, to take advantage of resource verification reduction. For more information about how to code the application capabilities vector, see [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

Use of the resource verification reduction function is limited to application programs using the VTAM LU 6.2 programming interface. In addition, the application program must use network-qualified names for session requests. Because some session requests require that the destination LU be contacted, verification reduction cannot be performed for sessions using:

- Application supplied dial parameters (ASDP)
- Cryptography
- Generic resources
- Secure Class of Service

- USERVARs

The benefit of resource verification reduction is limited for session requests traversing subnetwork boundaries (border nodes and interchange nodes). Resource verification reduction can eliminate a verification search across only one of the traversed subnetworks.

Controlling the use of resource verification reduction

The use of resource verification reduction is controlled by the VFYRED and VFYREDTI start options.

- Use of the VFYRED start option determines whether resource verification reduction is allowed for all LU 6.2 session initiation requests using the APPCCMD application programming interface.
- Use of the VFYREDTI start option indicates the maximum amount of time that cache information may be used before a resource location must be verified.

Resource verification reduction is suppressed for application programs with a LOGON exit unless the application program indicates support of multiple logon exit through the application capabilities control vector. Table 42 on page 332 shows when the resource verification reduction function is allowed and requested by the requesting application.

<i>Table 42. Resource verification reduction matrix</i>			
VFYRED start option value	LOGON exit included	Multiple logon indicator set on	Verification reduction allowed
YES	YES	YES	YES
YES	YES	NO	NO
YES	NO	N/A	YES
NO	N/A	N/A	NO

LU 6.2 session limits

VTAM maintains a set of session limits for partner logical units for each mode name used. These session limit values are used if a partner LU attempts to negotiate a change in the current number of sessions. If VTAM receives a request from the other partner LU to change the number of sessions for its LU 6.2 logical unit, VTAM uses the session limit values to determine whether to accept or reject the requested change. Note that the partner LU may be unable to accept increases in the session limits. Therefore, the end result of any change requests can still mean that the session limits do not increase.

If a session is available for use by a conversation, the process of starting and ending a conversation requires fewer system resources than starting and ending a session. Ideally, therefore, application programs should set up their session limits so that VTAM can establish a number of sessions with other LUs that will be relatively long-lived. Applications can then treat these sessions as resources that support conversations.

Requesting session limit changes: Application

An LU 6.2 logical unit requests that its partner LU alter the number of sessions for a given mode name by sending a change number of sessions (CNOS) request. A VTAM application program can request that its partner LU change the number of sessions. The application program does this by issuing an APPCCMD macroinstruction with the CONTROL=OPRCNTL and QUALIFY=CNOS operands specified.

Requesting session limit changes: Operator

You can enable the operator to modify CNOS parameters by specifying OPERCNOS=ALLOW on the APPL definition statement. The parameters can then be changed by a MODIFY CNOS command, which functions the same as an APPCCMD macroinstruction with the CONTROL=OPRCNTL and QUALIFY=CNOS operands specified. The operator can then also use the MODIFY DEFINE command, which functions the

same as an APPCCMD macroinstruction with the CONTROL=OPRCNTL and QUALIFY=DEFINE operands specified. The application is notified when the operator-issued commands have completed.

Coding session limits

These session limit values can be set for a VTAM application program by coding specific operands on the APPL definition statement. The VTAM application program can also establish these session limit values and override the specifications on the APPL definition statement by issuing an APPCCMD macroinstruction with the QUALIFY=DEFINE operand specified. Following are the unique session limit values that are used:

- The overall maximum number of sessions between two partner LUs for a given mode name, which you specify on the DSESLIM operand of the APPL definition statement
- The minimum number of sessions for a given mode name for which the application program is the contention winner, which you specify on the DMINWNL operand of the APPL definition statement
- The minimum number of sessions for a given mode name for which the application program might be the contention loser (the partner LU is the contention winner), which you specify on the DMINWNR operand of the APPL definition statement

The sum of the DMINWNL and DMINWNR values must always be less than or equal to the maximum number of sessions specified on the DSESLIM operand. The AUTOSSES operand on the APPL definition statement can be used to have VTAM automatically activate a specific number of contention-winner sessions for a given mode name for the application program. VTAM establishes these sessions when the first CNOS request is issued by the application program for a partner LU using a given mode name.

Session limits and deactivation

Using CNOS, an LU 6.2 session partner can request that the session limits be set to 0. Then, the sessions that are active and not used for a conversation are deactivated. The LU partner that is responsible for deactivating the sessions for a given mode name is determined by an indicator in the CNOS request. The DRESPL operand on the APPL definition statement can be used to indicate if the VTAM application program accepts the responsibility for deactivating the sessions (ALLOW) or if VTAM assigns that responsibility to the partner LU when the CNOS request is sent (NALLOW).

When sessions between partner LUs are to be deactivated, the existing conversations are not disrupted, but subsequent allocation requests for conversations can be queued. Before deactivating the sessions for a given mode name, the partner LUs can indicate whether conversations that are queued for a session are to be processed or discarded. The DDRAINL operand on the APPL definition statement indicates whether the VTAM application drains the queued allocation requests when deactivating sessions. If you code DDRAINL=ALLOW, the conversations that are queued for a session for the given mode name are processed before deactivating the sessions. If you allow conversations to be drained, new conversations are allocated a session if the allocation request occurs before the deactivation of all of the sessions for the given mode name. If you code DDRAINL=NALLOW on the APPL definition statement, VTAM does not allow the application program to drain queued session requests and the conversation requests are rejected.

Managing LU 6.2 sessions with operator commands

Session characteristics, such as session limits, LU capability, and LU responsibilities, must be managed for every conversation involving an LU 6.2 application program. These characteristics can be controlled by the application program and through operator commands.

Operator commands manage conversation resources by establishing or modifying:

- New session limits between an application and partner LU (MODIFY CNOS).
- Session limit values VTAM uses to negotiate CNOS requests initiated by a partner LU for a specific application (MODIFY DEFINE). These values are stored in the logon mode table.

You can establish session limits using the MODIFY CNOS operator command. After session limits are established, actual sessions can be started.

Session limits must also be established for a mode name before conversations that use that mode can be allocated. When a request for a session is made, the requested session limits can be accepted by the partner LU or the session limits can be changed to new values. The process for accepting or changing proposed session limits is called negotiation. Logon mode table values are used in session limits negotiation. Negotiation is handled for the application program by the MODIFY DEFINE operator command. For more information about session negotiation, see the [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

LU 6.2 security

VTAM LU 6.2 support provides the following security functions for your network:

- Encryption facility
- LU 6.2 conversation-level verification (user ID verification)
- LU 6.2 session-level LU-LU verification

The encryption facility protects data passing over lines between network resources by permitting enciphering and deciphering of data for LU-LU sessions. For an application program or peripheral node logical unit to have cryptographic sessions, the host processor must support cryptography. For more information about the encryption facility, see [“Cryptography facility” on page 289](#).

LU 6.2 user ID verification is a conversation-level security protocol, taking place at the time a conversation is started. For more information about LU 6.2 conversation-level security, see the [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

LU 6.2 session-level LU-LU verification is a session-level security protocol that is used to verify the identity of each logical unit at the time the session is activated. LU 6.2 session-level LU-LU verification provides the ability to verify the identity of an application program partner LUs during the activation of sessions between type 6.2 LUs. VTAM-generated random data is encrypted using one of the data encryption standard algorithms. The encryption key is a password associated with each LU-LU pair. The encrypted data is carried on both the session activation request and the response so that each LU partner can verify the other partner for the session. The passwords used for session-level LU-LU verification are not coded on any VTAM definition statements but are implemented through an external security management product, such as RACF. If you plan to use LU 6.2 session-level LU-LU verification, a security management product (such as RACF 1.9 or later) must be installed and active. In addition, a profile for the LU needs to be in the security management database. With RACF 1.9, the APPCLU class needs to be active and a profile of the LU needs to be in the APPCLU class. For an example of the appropriate RACF coding, see the [z/OS Security Server RACF Security Administrator's Guide](#).

During activation of LU 6.2 sessions involving control points, the VERIFYCP start option specifies whether VTAM performs session-level LU-LU verification.

The VERIFY and SECLVL operands on the APPL definition statement identify the level of partner-LU security verification.

If the application program is the PLU and an LU-LU password is defined for the partner LU, VTAM requests that LU-LU verification be performed during session activation. If a password is not defined, LU-LU verification is not requested.

If the application program is the secondary logical unit, one of the following conditions occurs:

- If the session activation request specifies LU-LU verification and the LU-LU password is defined, verification is performed.
- If the session activation request does not specify LU-LU verification and no LU-LU password is defined, session activation continues without verification.
- If neither of the previous cases applies, VTAM rejects the session activation request.

If you are using LU 6.2 session-level LU-LU verification, you must create the RACF profile using either three-part or four-part names. Create the profile using a four part name if one of the following is true.

- The profile you are creating is for the VTAM control point. That is, you are trying to verify the identity of the partner CP for CPCP sessions.

- The local LU supports network-qualified names. The local LU supports network-qualified names if it specified NQNames=YES on the OPEN ACB.

Create a RACF profile using a three part name if none of the above situations are true. See the following examples of a three part name definition followed by a four part name.

- RDEFINE APPCLU localnetid.localLU.remoteLU UACC(NONE) ...
- RDEFINE APPCLU localnetid.localLU.remoteLU UACC(NONE) ...

Note: Specifying VERIFY=OPTIONAL does not restrict the ability of a logical unit without a corresponding LU-LU profile to establish sessions with this application program.

If using RACF as your external security management product, the MODIFY PROFILES command enables you to reload an active application program set of existing defined RACF profiles. However, you cannot change the RACF profile with MODIFY PROFILES, only refresh it. The RACF profile contains the LU-LU password and only someone with RACF security clearance can change it. This can be helpful when the password for an LU-LU pair has been changed or when session activation errors are occurring. The profile changes affect only those sessions that are started after using the command; active sessions are not affected.

LU 6.2 sync point services

LU 6.2 sync point services is a session-level protocol. The maximum service level parameters are exchanged at session activation. However, when the conversation is allocated, the service level can be lower than the parameters exchanged at session activation. This can be caused by the inability of the other logical unit to support a higher service level. Synchronization of protected resources is the application program's responsibility. It is a means of recovery from errors or failures that allows the data network to divide its transmissions into logical units of work and commit or back out the information in these divisions of work.

The sync point function is divided between VTAM and a VTAM application program. LU 6.2 provides the facilities to allow the application program to implement a sync point facility by supporting the flow of information between application programs and the ability to control a session by rejecting, holding, or releasing a session based on a session identifier.

Persistent LU-LU session support does not allow LU 6.2 sync point sessions to be retained. Instead, they end when an application fails and they cannot be recovered.

The SYNCLVL operand on the APPL definition statement indicates the level of synchronization for this application program. If you do not code SYNCLVL, confirmation requests and sync point processing are not supported.

On the APPL definition statement, you can also specify (in conjunction with sync point support) when the ATTN exit for an LU 6.2 application program is to be scheduled. This exit routine can be invoked when the last session (ATNLOSS=LAST) is deactivated between an LU and MODE (LU-mode) pair. The ATTN exit can also be invoked for all session deactivations by specifying a value of ALL on the ATNLOSS operand. The LU 6.2 application program can use this operand without implementing sync point support.

For more details on sync point support, see [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

Selective termination of idle LU 6.2 sessions

You can limit the use of some network connections, such as lines. Defining a limited network connection enables you to limit the use of the connection and keep only active sessions on that connection. Sessions that are not active are terminated. If all sessions are terminated, then the connection is terminated. The best network connections to define as limited are lines and PUs whose cost is determined by the length of time a connection exists. Defining these as limited can help reduce switched line connect charges.

When a session traverses a limited connection, the session can be terminated if no conversation is detected for a set period of time. You must specify the timeout on the LIMQSINT operand of the APPL

definition statement. LIMQSINT specifies how long (in seconds) unused limited resource sessions will remain on the queue before they are terminated.

Note: Only LU 6.2 conversations are affected by limited resource definition. Non-LU 6.2 sessions are unaffected and cannot be limited.

To use LU 6.2 limited resource management:

1. Determine which network connections you want to define as limited.
2. Define the network connections to VTAM by coding LIMRES=YES on the GROUP, LINE, or PU statement. LIMRES is valid for PU types 1, 2, and 2.1 only.
3. Determine the time interval for VTAM to search the queue and define the interval to VTAM using the LIMQSINT operand on the APPL definition statement.

You can specify a line, a group of lines, or a PU as a limited resource for the following major nodes:

- External communication adapter
- Local SNA (PU only)
- MODEL (PU only)
- NCP
- Switched (PU only)

After you define the resource to be limited, set the time interval for the queue search. As a starting point, you should specify a time interval at least 1 second less than half of the shortest line time cost interval. For example, if the line time cost interval is 120 seconds, you should specify 59 seconds on the LIMQSINT operand. You might need to adjust the time interval for different applications.

All idle associations are ended. Associations that were specifically requested by an application program (associations for which an application program issued an ACF.Associate request) are not affected by selective termination at the system where the ACF.Associate request was initiated, but can be terminated automatically on the system that was the target of the association.

Selective termination of idle network management sessions

Network management is the process of planning, organizing, monitoring, and controlling a communication-oriented data processing or information system. The architecture provided to assist in network management of SNA systems is called *management services* and is implemented as a set of functions and services designed to capture and use the information needed for effective management.

VTAM can exchange management services information with other systems that provide management services transport, or multiple-domain support, such as AS/400, Communication Manager/2, or other VTAM systems. This type of information flows between architected transaction programs in the respective systems over SNA sessions. The network management sessions that carry this exchange use either the CPSVCMG or SNASVCMG mode names. For further information about management services concepts and architecture, see *Systems Network Architecture Management Services Reference*.

The limited resource function is also available on some control point sessions used for network management flows, though for these sessions, it differs from the way it operates for noncontrol point sessions. For network management flows using the CP SNASVCMG session, a session is deactivated only when there are not any management services (MS) transport transactions outstanding, in addition to the session being inactive.

By coding the LIMINTCP start option, you specify the interval to retain a free CP SNASVCMG session for a resource that is defined as limited (with the LIMRES operand). The CP SNASVCMG session is used for some network management flows.

Note: The limited resource function for the control point is supported only on CP sessions for network management flows using the CP SNASVCMG session. The function is not available on APPN CP sessions or dependent LU server sessions.

To define limited resource management for network management flows, perform the following steps:

1. Determine which network connections you want to define as limited. These are the lines whose costs are determined by the length of connect time.
2. Define the network connections to VTAM by coding LIMRES=YES on the GROUP, LINE, or PU definition statement. LIMRES is valid only for PU type 2.1.

You can specify a line, a group of lines, or a PU as a limited resource for the following major nodes:

- External communication adapter
- Local SNA (PU only)
- MODEL (PU only)
- NCP
- Switched (PU only)

3. Determine the interval for VTAM to search the queue.

As a starting point, the interval should be at least 1 second less than half the shortest line time cost interval. For example, if the shortest time cost interval is 1 minute, the interval should be 29 seconds.

Specify the interval to VTAM with the LIMINTCP start option. If you do not code LIMINTCP, the limited resource function will not be supported for network management flows. See [z/OS Communications Server: SNA Resource Definition Reference](#) for information about how to specify the LIMINTCP start option.

APING support

APING support provides the ability to test for connectivity between LU 6.2 resources, and to display pertinent routing and performance information. APING support consists of two transaction programs, APING and APINGD. The APING transaction program represents the client side of a transaction, while the APINGD transaction program represents the server side of a transaction. When a conversation is established, the client transaction program sends packets of data to the server transaction program. When the client is finished, it indicates to the server transaction program either to send the data back to the client or to send a confirmation back to the client indicating that the data has been received. The client continues the transaction until complete.

The VTAM APING transaction program is started when a DISPLAY APING command is issued. The APINGD transaction program is started when another system in the network attempts to verify its connectivity and performance with VTAM using the APING transaction program.

Using the DISPLAY APING command, the network operator can do the following things:

- Verify connectivity with any LU 6.2 resource in the network
- Verify that another VTAM 4.3 or later node is operational
- Check the performance of the network using a particular logon mode
- Display routing information to the destination node, if a new session is established for the APING transaction

When you issue a DISPLAY APING command, the following operands determine the amount of data that is sent by the client transaction program to the server transaction program:

SIZE

Specifies the size (in bytes) of the packets to be sent.

CONSEC

Specifies the number of consecutive packets to be sent.

Note: Depending on the maximum allowed RU size on the session path, the packets might be transmitted in a number of flows different than what you specified for CONSEC.

ITER

Specifies the number of times that consecutive packets are sent to the server transaction program and returned.

ECHO

Specifies whether data is returned to the client transaction program.

For example, if you issue the command

D NET,APING,ID=d1uname,CONSEC=3,ECHO=YES,ITER=2,SIZE=500, as shown in [Figure 98](#) on [page 338](#), three consecutive packets of 500 bytes are sent to the server transaction program two times, and the server transaction program returns three packets of 500 bytes two times.




Figure 98. Example of flows between client and server for DISPLAY APING command

Guideline: The maximum value for CONSEC, ITER, and SIZE is 32763. By coding 32763 for each of these operands with ECHO=YES, 32763 packets of 32763 bytes of data are sent to the partner transaction program 32763 times and the partner transaction program returns 32763 packets of 32763 bytes of data 32763 times. This could cause a significant impact to the performance of your network.

The number of concurrent APING or APINGD transaction program instances allowed can be displayed with the DISPLAY APINGTP and DISPLAY APINGDTP operator commands, and can be adjusted using the MODIFY APINGTP and MODIFY APINGDTP commands. The default instance limit for the APING and APINGD transaction programs is 10. If you need to cancel APING transaction processing, see [“Canceling APING sessions”](#) on [page 338](#). The command verification exit, ISTCMMND, can be used to restrict the operands specified on the DISPLAY APING command.

Canceling APING sessions

If you want to terminate the sessions used by APING transaction processing, you can issue a VARY TERM command with the session identifier (SID) of the session being used for the data.

Canceling APING processing

If a session is allocated for the DISPLAY APING command, the SID of the session is reported in message IST1490I of the routing information message group, or as output from a DISPLAY APINGTP,LIST=ALL command. If a session is not allocated for the DISPLAY APING command, the routing information message group is not issued. The performance information message group is not issued until the transaction processing is complete.

Procedure

You can determine which session to cancel in the following way:

1. Issue a D NET,APINGTP,LIST=ALL command.
2. Find the session with the LUNAME specified on the DISPLAY APING command.
3. Note the session identifier.
4. Issue a V NET,TERM,SID=session_id command to terminate the session.

Canceling APINGD processing

If this VTAM is the server for an APING transaction, you need to cancel APINGD processing to terminate it.

Procedure

You can determine which session to cancel in the following way:

1. Issue a D NET,APINGDTP,LIST=ALL command.
2. Note the session identifiers.
3. Issue a V NET,TERM,SID=session_id command for each session you want to terminate.

High-performance data transfer (HPDT)

High-performance data transfer (HPDT) optimizes the performance of large data transfers for VTAM LU 6.2 applications. The performance benefits provided by HPDT services are achieved by more efficient use of critical CPU resources (CPU cycles, cache, memory bus, and the channel). Specifically, system resources are conserved by the following methods:

Reduction of VTAM internal data copies:

The movement of large pieces of data can result in significant processing costs on the send and receive execution paths. This cost is because of the effective path length of the MVCL instruction. Additionally, data movement flushes the CPU cache and consumes the memory bus bandwidth.

Reduced path length:

The processing of data transfers is streamlined through the use of the communications storage manager (CSM), which reduces the increase in path length that normally occurs as API crossing and PIU size increase. CSM is described in [Appendix C, “Communications storage manager,” on page 581](#).

Use of HPDT MPC:

HPDT services take advantage of HPDT MPC, which is described in [“Multipath channel connections” on page 40](#).

HPDT includes a service that reduces the number of times data is moved for an APPCCMD send or receive request. No application change is required to receive performance benefits. However, an HPDT interface is also provided to enable applications to obtain performance benefits even greater than those provided by HPDT services. For more information about the HPDT interface, see [z/OS Communications Server: SNA Programmer's LU 6.2 Guide](#).

Applications using the HPDT interface can specify CSM data space or CSM extended common service area (ECSA) for inbound data. When CSM data space is specified, the storage, when fixed, tends to be backed by real storage on or above the 2-gigabyte bar. VTAM receives data over the channel in the storage type specified by the STORAGE operand on the TRLE definition statement. If a mismatch exists between the VTAM TRLE definition and the application specification, an extra data move occurs. You should consider the CSM storage type used by your LU 6.2 applications when setting the STORAGE parameter for the

TRLE definition. For more information about the TRLE definition, see [z/OS Communications Server: SNA Resource Definition Reference](#).

Communications storage manager (CSM)

Using the communications storage manager (CSM) function, authorized host applications can share data with other CSM users without having to physically move the data. A CSM user can be any system-authorized application program or product. Users of CSM obtain and return storage in the form of CSM buffer pools by using the IVTCSM macroinstruction.

VTAM uses CSM to perform channel I/O over a HPDT MPC connection and to provide the high-performance data transfer (HPDT) function for host LU 6.2 applications. (HPDT MPC connections are explained in [“Multipath channel connections”](#) on page 40.) Host LU 6.2 applications use the IVTCSM and APPCCMD macroinstructions to reduce the use of system resources for large data transfers. For more information about the HPDT function, see [“High-performance data transfer \(HPDT\)”](#) on page 339.

The application programming interface for CSM is described in [z/OS Communications Server: CSM Guide](#).

For information about how to install, configure, monitor, and diagnose CSM storage, see [Appendix C, “Communications storage manager,”](#) on page 581.

Chapter 14. Functions provided by VTAM in a sysplex

A sysplex is a set of MVS systems communicating and cooperating with each other, through certain multisystem hardware components and software services, to process customer workloads. The inclusion of a coupling facility within the sysplex allows for high performance data sharing. For more information about sysplex, see [z/OS MVS Setting Up a Sysplex](#).

VTAM supports the attachment of a sysplex to a network. Sessions into the sysplex can be established from either subarea nodes or APPN nodes. There are several VTAM functions that are available only in a sysplex environment. If a coupling facility and an APPN or mixed APPN and subarea environment exists in the sysplex, the user can take advantage of the following VTAM functions:

Generic resources

Allows an user to easily connect to any one of a number of duplicate application programs on different systems. The user uses a generic name and VTAM determines the actual application program for the session based on workload and other performance criteria.

Multinode persistent sessions

Allows for the restoring of an application program session during a planned or forced takeover, or after hardware, operating system, or VTAM failures.

TSO generic resources

Allows for multiple TSO/VTAM application programs to have a common name. During session establishment, VTAM determines the actual application program to be used for the session.

Dynamic definition of VTAM-to-VTAM connections

Allows for dynamic establishment of XCF connections between VTAM nodes in a sysplex when VTAM is initialized with the XCFINIT start parameter, or at a later time when the operator determines connections can be established (by activating the XCF Local SNA major node, ISTLSXCF).

VTAM also provides support for TCP/IP functions that need access to a coupling facility. If a coupling facility exists in the sysplex, the user can take advantage of the following TCP/IP functions:

Sysplexports

Enables multiple TCP/IP stacks in the sysplex to collaborate on the assignment of ephemeral port numbers for a distributed dynamic virtual IP Address (DRVIPA) that is shared by a number of applications. This function uses a structure in the coupling facility to track the ephemeral port assignments across the sysplex.

Sysplex-wide security associations

Allows for IPsec security associations (SA) and their workloads to be distributed to target TCP/IP stacks within the sysplex. It also allows the SAs to be automatically restarted on another processor in the sysplex when a dynamic routable virtual IP Address (DRVIPA) takeover occurs. This takeover capability is provided for the following cases:

- Takeover occurs because the VIPA owning stack or its host went down (unplanned takeover)
- Takeover or giveback occurs when the stack and its host stay up but a VIPA ownership change between hosts is required (planned takeover)

Setting up the sysplex environment for VTAM and TCP/IP functions

The following sections apply to VTAM functions.

Sysplex subplexing

VTAM participates in the sysplex environment through two sysplex groups:

- The XCF group is used for dynamic definition of VTAM-to-VTAM connections.
- The CFS group is used for VTAM coupling facility services.

You partition VTAMs in the sysplex into subsets (subplexes) by modifying the names of the two sysplex groups that the sysplex joins. Use subplexing to separate and isolate the XCF connectivity of sets of VTAM nodes in the sysplex. VTAMs using the same XCF group name and the same CFS group name are in one subplex, and VTAMs with a different XCF group name and CFS group name are in a different subplex.

Modify the group names that VTAM uses with the XCFGRPID start option. VTAM joins the XCF group with a name in the format ISTXCFvv and it joins the CFS group with a name in the format ISTCFSvv, where vv is the numeric 2-digit value you supply on the XCFGRPID start option. The value specified for vv must be in the range 02-31. If you specify a single digit in the range 2-9, the value is padded on the left with 0. If you do not supply an XCFGRPID value, the XCF group name is ISTXCF and the CFS group name is ISTCFS01.

TCP/IP stacks can also be partitioned into subsets (subplexes) within the sysplex by specifying the XCFGRPID parameter on the GLOBALCONFIG statement in the TCP/IP profile. Because the TCP/IP stack relies on VTAM support for its XCF connectivity, TCP/IP subplexes cannot span VTAM subplexes. That is, two TCP/IP stacks in the same subplex cannot be in different VTAM subplexes. To enforce this, the sysplex group that TCP/IP joins in the sysplex has a name in the format EZBTvvtt, where vv is the VTAM group ID suffix you specified on the XCFGRPID start option, and tt is the TCP/IP group ID suffix value you specified on the XCFGRPID parameter on the GLOBALCONFIG statement. For example, if VTAM is started with a value of XCFGRPID=11, and a TCP/IP stack is started with a GLOBALCONFIG statement specifying XCFGRPID=21, the corresponding TCP/IP sysplex group name is EZBT1121.

If VTAM is stopped and restarted with a different value for the XCFGRPID parameter, the TCP/IP stacks that use that VTAM must also be stopped and restarted to pick up the new VTAM subplex suffix.

Because subplexing separates XCF connectivity, the VTAM and TCP/IP coupling facility structures that are accessed must be separated by subplexing.

- For the VTAM structures (for generic resources and MNPS), VTAM appends the VTAM XCFGRPID suffix to the end of the name specified on the STRGR and STRMNPS start options (or on ISTGENERIC and ISTMNPS, if the names are the default values). Thus, using the previous example, if you specify STRGR=ISTMYGR and STRMNPS=ISTMYMNPS as start options, and you also specify XCFGRPID=11, VTAM would attempt to connect to the structures ISTMYGR11 and ISTMYMNPS11. The fully suffixed MNPS and generic resources structure names can be displayed by issuing the D NET,ID=VTAM command.
- For the TCP/IP structures (for SWSA and sysplexports), VTAM appends the VTAM XCFGRPID suffix, followed by the TCP/IP XCFGRPID suffix to the end of the base structure names to create the separation. The base structure name for SWSA is EZBDVIPA and the base structure name for sysplexports is EZBEPOR. Again, using the previous example, TCP establishes connectivity to structures EZBDVIPA1121 and EZBEPOR1121.

Note: Any system in the sysplex that has access to the coupling facilities can be picked as a monitor for sysplex processing, and this has operational implications affecting how certain messages are processed. It is possible that the XCF monitor processing could be issuing messages on a system in 'subplex a' for a system that resides in 'subplex b'.

References to the base group names, ISTXCF, ISTCFS01, EZBTCPCS, and the base structure names, ISTGENERIC, ISTMNPS, EZBDVIPA, and EZBEPOR reference the fully suffixed names, if subplexing is used.

Setting up a subplex

Setting up TCP/IP and VTAM subplexes within a sysplex requires some preparation.

Procedure

To prepare your sysplex for subplexing, perform the following steps:

1. If there are TCP/IP stacks that use the subplexed VTAM nodes, see [setting up a sysplex in topic in z/OS Communications Server: IP Configuration Guide](#).
2. If the TCP/IP stack subplex configuration (determined in the previous step) requires the use of HyperSockets VLAN IDs (using the IQDVLANID parameter on the GLOBALCONFIG TCP/IP profile statement), ensure that all VTAM nodes in the sysplex are at z/OS V1R8 or later.

3. If VTAM coupling facility structures for generic resources or MNPS are to be used, ensure that all the structures have been defined to z/OS and are specified in the active CFRM policy using their fully suffixed names. For example, if the generic resources function is to be used with the base structure name of ISTGENERIC, and two VTAM subplexes are to be configured, one subplex with a suffix of 11 and the other with a suffix of 12, then ensure that the ISTGENERIC11 and ISTGENERIC12 structures are defined to MVS and specified in the active CFRM policy.
4. For each VTAM that needs to specify an XCFGRPID value or supports TCP/IP stacks that need to be changed to specify an XCFGRPID value or IQDVLANID value on their GLOBALCONFIG statements, do the following steps:
 - a. If this VTAM requires a new XCFGRPID value, stop VTAM
 - b. For each TCP/IP stack supported by this VTAM that requires a new XCFGRPID or IQDVLANID value:
 - i) Stop the TCP/IP stack.
 - ii) Update the GLOBALCONFIG statement in the TCP/IP profile for this stack with the new XCFGRPID and IQDVLANID values.
 - iii) Restart the TCP/IP stack.
 - c. If VTAM was stopped, update the VTAM start options to add the new XCFGRPID option with the group ID value for the new subplex (or specify the new XCFGRPID start option when prompted, after starting VTAM).
 - d. Restart the VTAM node.

Considerations

Both the XCF and CFS groups must be considered when planning sysplex coupled data sets. Specifically for the XCF group of VTAM, the MAXMSG parameter for PATHIN and PATHOUT should be adequately tuned to minimize the number of no buffer conditions. Inadequate tuning of MAXMSG can lead to unexpected XCF outages (INOPs). See [z/OS MVS Setting Up a Sysplex](#), for additional information.

The following information applies to all VTAM and TCP/IP functions that require a coupling facility structure in an MVS sysplex. If you plan to use the generic resource function, the multinode persistent session function, the sysplexports function, or the Sysplex Wide Security Associations function, a coupling facility structure must exist for each function.

- An MVS coupling facility provides the hardware and software that supports high-speed shared storage across multiple MVS systems and provides multiple copies or images of the MVS operating system to process work concurrently.
- There must be an active coupling facility resource management (CFRM) policy for the sysplex environment.

A CFRM policy defines how an installation will manage its coupling facility resources, which can be shared among multiple subsystems. For more information about how to create, activate, and make changes to the CFRM policy, see [z/OS MVS Setting Up a Sysplex](#).

The CFRM policy information describes the coupling facilities that are used in the sysplex and specifies the requirements for structure size and allocation within each coupling facility.

- Each coupling facility structure to be used must be defined in the active CFRM policy for the sysplex. If subplexing is used, the structure names specified in the CFRM policy must include the subplex suffixes. You might need to define (using the CFRM policy) more than one structure of each type as being in the sysplex. See [“Sample CFRM coding” on page 346](#) for an example of CFRM coding.
- When defining the coupling facility structure in a CFRM policy, specify the amount of coupling facility storage which should be allocated on behalf of this structure. For information about determining the amount of storage necessary, see [“Determining the size of the coupling facility structure” on page 346](#).
- All VTAM and TCP/IP functions using a coupling facility are designed to maintain data integrity when any single type of failure occurs, but do not reliably handle concurrent failures of different types. For instance, if a VTAM node and the coupling facility structure fail concurrently, data is lost. Thus, it is

advisable to isolate the coupling facilities from the VTAM nodes (for example, they should not share the same power supply).

In the same way, multiple coupling facilities in the same configuration should be isolated from each other. If more than one coupling facility is present in a configuration, a new version of the structure can be created if the coupling facility containing the active version of the structure fails. Also, if a link failure to the structure occurs, the structure might be moved to a new coupling facility where it could be accessed by all VTAM nodes in the configuration.

Coupling facility structure attributes

The attributes defined by VTAM describe the various coupling facility structure characteristics. For the structures used by VTAM and TCP/IP functions, VTAM supplies all attributes to MVS listed in [Table 43 on page 344](#) and [Table 44 on page 345](#) with the exception of the maximum list counts attribute. The list-set-entry count or list-set-element count attribute is used to determine the size specified in the policy.

When the structure size has been determined, the corresponding CFRM policy can be defined following the directions in [z/OS MVS Setting Up a Sysplex](#). For the sysplex to recognize a structure, a CFRM policy must be defined and active. This definition indicates that the coupling facility will contain the named structure.

Access to the structure occurs when VTAM connects to a coupling facility structure.

<i>Table 43. Structure attributes used to compute structure size</i>				
Attribute	Generic resources	Multinode persistent sessions	Sysplex Wide Security Associations	Sysplexports
Adjunct	Yes	Yes	Yes	Yes
List entry reference	Names	Keys	Keys	Keys
List element characteristic	2	0	0	0
List count	4	256 ¹	See DVLSTCNT start option	1024
Number of lock entries	3	0	0	1024
Maximum number of data elements	16	255	255	32
Initial entry-to-element ratio ²	List entry portion – 50 Data object portion – 1	List entry portion – 1 Data object portion – 2	List entry portion – 1 Data object portion – 10	List entry portion – 1 Data object portion – 32
Maximum list set element count	N/A	See “Determining the size of the coupling facility structure” on page 346	N/A	N/A
Maximum list set entry count	See “Determining the size of the coupling facility structure” on page 346	N/A	See “Determining the size of the coupling facility structure” on page 346	See “Determining the size of the coupling facility structure” on page 346

Table 43. Structure attributes used to compute structure size (continued)

Attribute	Generic resources	Multinode persistent sessions	Sysplex Wide Security Associations	Sysplexports
Lock table entry characteristic	0	0	0	0
Notes: <ol style="list-style-type: none"> 1. The list count number can increase across a structure rebuild. Issue DISPLAY STATS,TYPE=CFS to get the current number of lists allocated. 2. During the rebuild process this ratio might be altered to better fit the actual use of list entries and elements. VTAM can also use the dynamic structure alter to change this ratio. 				

Table 44 on page 345 contains structure attributes used to determine the coupling facility preference and exclusion list for a structure.

Table 44. Structure attributes defined by VTAM

Attribute	Generic resources	Multinode persistent sessions	Sysplex Wide Security Associations	Sysplexports
Structure type	List	List	List	List
Structure rebuilds allowed	Yes	Yes	Yes	Yes
Structure connection disposition	KEEP	KEEP	DELETE	DELETE
Structure disposition	DELETE	KEEP	DELETE	DELETE

Structure type

The coupling facility supports three structure types: cache, list, and lock. The coupling facility structure type used by VTAM is list. List structures allow list format data sharing. The type attribute is defined when the first VTAM node connects to the coupling facility structure.

Structure rebuilds

Structure rebuild allows a connector to a structure to allocate another structure with the same name and reconstruct data in the newly allocated structure. Rebuilding allows connectors to change the location or attributes of the structure without having to disconnect.

VTAM sets the size of the rebuilt structure equal to the current size of the original structure, unless there is a structure storage shortage or the policy size information is being changed across the rebuild. VTAM considers a shortage to exist if more than 80 percent of total structure storage is in use. If there is a shortage, and the current size is less than the maximum size, the rebuild structure size will be set to the maximum size. If the CFRM policy size information is being changed across the rebuild, the rebuild structure size will be set as follows:

- If INITSIZE is not specified, the rebuild structure size is the value of SIZE in the CFRM policy definition.
- If INITSIZE is specified and:
 - The current structure size is greater than or equal to the INITSIZE value and less than or equal to the SIZE value, the current size is retained.

- The current structure size is less than the INITSIZE value, the INITSIZE value is used.
- The current structure size is greater than the SIZE value, the SIZE value is used.

See [“Structure rebuild” on page 348](#) for more information about structure rebuilds.

Structure connection disposition

The structure connection disposition attribute indicates how the connection is handled in the event of abnormal disconnects from the structure.

DELETE

Indicates that if a VTAM node fails or is forced to disconnect because of some type of failure, the node's connection definition is removed.

KEEP

Indicates that if a VTAM node fails or is forced to disconnect because of some type of failure, the node's connection remains defined but is in a failed-persistent state. Failed-persistent state indicates that if a VTAM node fails or is forced to disconnect because of some type of failure, the node's connection remains defined but is in a failed-persistent state. The failed-persistent state also indicates that recoverable data is located in the coupling facility structure.

Note: A network operator might see VTAM in a failed-persistent state because of a VTAM failure or because VTAM disconnects from the coupling facility. To ensure that data needed to recover is available, the operator should not issue the MVS SETXCF FORCE command.

Structure disposition

Structure disposition indicates what happens to the structure when there are no defined connections to a structure.

DELETE

Indicates that the structure will be deallocated after the last connection to that structure is deleted.

KEEP

Indicates that the structure is never deallocated. The only way to deallocate the structure is through the MVS SETXCF FORCE command.

Determining the size of the coupling facility structure

The Coupling Facility Structure Sizer Tool (CFSizer) is used to determine the following values:

- The values for the INITSIZE and SIZE parameters that are specified in the CFRM policy for a VTAM or TCP/IP structure.
- The value for the DVLSTCNT start option for an EZBDVIPA structure.

The CFSizer tool can be found at the following web address: <http://www.ibm.com/systems/support/z/cfsizer/>.

The CFSizer tool should be used to determine the INITSIZE, SIZE, and DVLSTCNT values whenever a new CFLEVEL of the Coupling Facility Control Code has been loaded into a Coupling Facility.

Sample CFRM coding

The following example shows the CFRM policy definition for including the generic resources, MNPS, SWSA and sysplexports structures in a sysplex that is partitioned into two VTAM subplexes, 11 and 12. Within VTAM subplex 11, the sysplex is partitioned into two TCP/IP subplexes, 21 and 22. Within VTAM subplex 12, the sysplex is partitioned into two TCP/IP subplexes, 22 and 23. Note that TCP/IP subplex 22 within VTAM subplex 11 is separate from TCP/IP subplex 22 within VTAM subplex 12. TCP/IP subplexes cannot span VTAM subplexes.

```
//SYSIN DD *
        DATA TYPE(CFRM) REPORT(YES)
        DEFINE POLICY NAME(VTAMSAMP) REPLACE(YES)
```



```

CF NAME(RALNSCF1)
  TYPE(009672)
  MFG(IBM)
  PLANT(02)
  SEQUENCE(000000041266)
  PARTITION(1)
  CPCID(00)
  DUMPSPACE(2000)

CF NAME(RALNSCF2)
  TYPE(009672)
  MFG(IBM)
  PLANT(02)
  SEQUENCE(000000041266)
  PARTITION(2)
  CPCID(00)
  DUMPSPACE(2000)

STRUCTURE NAME(ISTGENERIC11)
  SIZE(10000)
  INITSIZE(7000)
  PREFLIST(RALNSCF1,RALNSCF2)

STRUCTURE NAME(ISTGENERIC12)
  SIZE(10000)
  INITSIZE(7000)
  PREFLIST(RALNSCF1,RALNSCF2)

STRUCTURE NAME(ISTMNPS11)
  SIZE(25600)
  INITSIZE(12288)
  REBUILDPERCENT(30)
  PREFLIST(RALNSCF1,RALNSCF2)

STRUCTURE NAME(ISTMNPS12)
  SIZE(25600)
  INITSIZE(12288)
  REBUILDPERCENT(30)
  PREFLIST(RALNSCF1,RALNSCF2)

STRUCTURE NAME(EZBDVIPA1121)
  SIZE(50000)
  INITSIZE(15000)
  REBUILDPERCENT(20)
  PREFLIST(RALNSCF2,RALNSCF1)

STRUCTURE NAME(EZBDVIPA1122)
  SIZE(50000)
  INITSIZE(15000)
  REBUILDPERCENT(20)
  PREFLIST(RALNSCF2,RALNSCF1)

STRUCTURE NAME(EZBDVIPA1222)
  SIZE(50000)
  INITSIZE(15000)
  REBUILDPERCENT(20)
  PREFLIST(RALNSCF2,RALNSCF1)

STRUCTURE NAME(EZBDVIPA1223)
  SIZE(50000)
  INITSIZE(15000)
  REBUILDPERCENT(20)
  PREFLIST(RALNSCF2,RALNSCF1)

STRUCTURE NAME(EZBEPOR1121)
  SIZE(10000)
  INITSIZE(5000)
  REBUILDPERCENT(20)
  PREFLIST(RALNSCF2,RALNSCF1)

STRUCTURE NAME(EZBEPOR1122)
  SIZE(10000)
  INITSIZE(5000)
  REBUILDPERCENT(20)
  PREFLIST(RALNSCF2,RALNSCF1)

STRUCTURE NAME(EZBEPOR1222)
  SIZE(10000)

```

```

INITSIZE(5000)
REBUILDPERCENT(20)
PREFLIST(RALNSCF2,RALNSCF1)

STRUCTURE NAME(EZBEP0RT1223)
SIZE(10000)
INITSIZE(5000)
REBUILDPERCENT(20)
PREFLIST(RALNSCF2,RALNSCF1)

```

Connecting to and allocating storage for coupling facility structures

When a coupling facility resource management (CFRM) policy with a VTAM structure is activated, VTAM is notified and connects to the structure. A structure policy can be changed or activated before or after VTAM is active. VTAM attempts to connect to a VTAM coupling facility structure when it detects that a structure has been defined. For the TCP/IP functions, VTAM does not attempt to connect to the structure associated with that function until a TCP/IP stack requests access to the structure. Before a VTAM node attempts to connect to the coupling facility structure, it first checks to see that it is running on a version of MVS that supports the coupling facility and that the coupling facility structure is defined to the active CFRM policy. The first VTAM node to connect to a structure causes the allocation of storage for the structure within a coupling facility and also defines the structure attributes associated with the structure.

Structure rebuild

The structure rebuild process is used to create a new copy of a coupling facility structure, either in the same coupling facility or in a different coupling facility, and to store in this new copy of the structure all of the data from the old copy of the structure.

Notes:

1. For generic resources only, data is copied one of two ways during the structure rebuild process for generic resource structure. Data is either copied by a VTAM node in the generic resource configuration from the old version of the structure into the new version of the structure, or all VTAM nodes in the generic resource configuration each supply the new version of the structure with local data.
2. For multinode persistent sessions only, use the MVS FORCE command specifying the ARM parameter to cancel VTAM when a rebuild is in progress. This allows VTAM to correctly disconnect from the structure.

Initiating a structure rebuild

Note: VTAM can specify a connection specific reason for starting a rebuild. This reason is displayed in MVS message IXC526I. See [z/OS Communications Server: IP and SNA Codes](#) for a list of VTAM connection-specific reasons for starting a rebuild.

There are several reasons why a structure rebuild might be initiated:

- An MVS operator can initiate a structure rebuild
- VTAM will honor the REBUILDPERCENT value coded in the CFRM policy, if there is an active SFM policy. When a SFM is active, VTAM will allow MVS to initiate the rebuild based on the REBUILDPERCENT value. If no SFM policy is active, VTAM will initiate a structure rebuild when it loses access to the structure. For more information about the REBUILDPERCENT parameter and the SFM policy, see [z/OS MVS Setting Up a Sysplex](#). MVS can initiate a rebuild if a VTAM node loses access to the structure and the REBUILDPERCENT parameter is coded in the active CFRM policy and an SFM policy is active. For more information about the REBUILDPERCENT parameter and the SFM policy, see the [z/OS MVS Setting Up a Sysplex](#).
- A VTAM node can initiate a structure rebuild when notified that the structure has failed.

Stopping a structure rebuild

Note: VTAM can specify a connection-specific reason for stopping a rebuild. This reason is displayed in MVS message IXC527I. See the [z/OS Communications Server: IP and SNA Codes](#) for a list of VTAM connection-specific reasons for starting a rebuild.

When a structure rebuild has been initiated, it can be stopped for any one of several reasons.

- An MVS operator can stop a structure rebuild for structure.
- MVS automatically stops a rebuild if the new structure being created fails or becomes inaccessible.
- A VTAM node in a generic resource configuration can stop a structure rebuild if it loses access to the new version of the structure. A VTAM node never stops a structure rebuild when the old version of the structure has failed.
- A VTAM node in a generic resource configuration can stop a structure rebuild when it determines that all of the data in the old version of the structure might not be available for creating the new version of the structure.

When a structure rebuild is stopped, the old version of the structure continues to be used (if possible) and the new version of the structure, if allocated, will be deallocated.

If MVS or an operator initiates the rebuild, VTAM will not stop the rebuild.

Coupling facility duplexing

All z/OS Communications Server coupling facility structures support system-managed duplexing. This allows two copies of the structure to be created for backup purposes. Updates to the structure are automatically kept in sync so if access to one of the structures is lost, the other copy is used without any disruptions. To create a duplexed copy of the structure, use the SETXCF START,REBUILD,DUPLEX command. To stop duplexing and return to a single allocated structure, use the SETXCF STOP,REBUILD,DUPLEX command. For more information about the SETXCF command, see [z/OS MVS System Commands](#).

Coupling facility storage shortages

VTAM monitors the total amount of storage used by a coupling facility structure, the number of entries in use, and the number of elements in use. If VTAM detects storage 80 percent used (or greater), VTAM will issue message IST1439I, which indicates which type of storage shortage has occurred. The shortage situation is not ended until the constrained resource reaches 70 percent utilization.

To correct constrained resources:

- For total storage problems, change the SIZE parameter on the structure definition in the CFRM policy and issue a structure rebuild.
- For entries and element problems, VTAM performs a structure rebuild to adjust the entry-to-element ratio at 100 percent utilization.

Dynamic altering of structures

If you are using Level 1 Coupling Facility, VTAM provides nondisruptive reconfiguration of the coupling facility size and entry-to-element ratio attributes when VTAM detects a coupling facility storage problem.

Note: The MVS operator can use the MVS SETXCF ALTER command to initiate a structure alter process.

Dynamic changes to coupling facility size

For VTAM to dynamically increase the coupling facility storage size, specify the following on the structure policy statement:

SIZE

Specify the maximum allowed size for the structure.

INITSIZE

Specify the initial size of the structure. See [“Determining the size of the coupling facility structure”](#) on page 346 for information about calculating the structure size.

Note: For VTAM to dynamically alter the structure size, the SIZE value must be larger than the value of INITSIZE. The structure alter process is initiated to change the structure size when the current in-use percentage of total storage reaches 90 percent and the maximum size is greater than the current size. When the size has reached the value of SIZE attribute, the policy must be changed and an alter or rebuild command issued to incorporate the changes. Set the ALLOWAUTOALT parameter on the CFRM policy for VTAM or TCP/IP structures to ALLOWAUTOALT(NO) (or allow the ALLOWAUTOALT parameter to default to NO) to prevent MVS attempts to resize the structure from conflicting with VTAM re-sizing events.

Dynamic changes to entry-to-element ratio

The structure alter process is initiated to change the entry-to-element ratio when the current in-use percentage of either entries or elements reaches 90 percent.

Dynamic definition of VTAM-to-VTAM connections

Dynamic definition of VTAM-to-VTAM connections depends on the functions of XCF in a sysplex. The MVS system programmer must establish the XCF environment. VTAM requires that a value for the MVS symbol &SYSCONE be provided during MVS initialization. VTAM joins the XCF group named ISTXCF. No action is required to specify this group.

VTAM uses the XCF facilities only if they exist. If XCF facilities are not available or if subarea routing is used exclusively, VTAM-VTAM communication within a sysplex requires explicit definition.

Whether VTAM attempts to establish connectivity through XCF with other nodes in the sysplex during VTAM initialization is determined by the value specified for the start option XCFINIT.

- If XCFINIT=YES is specified (for an APPN node), VTAM joins the ISTXCF group at VTAM initialization. As it becomes aware of each other member node of the Sysplex, VTAM builds a TRLE and APPN PU definition for that node and begins the activation of the APPN PU. XCFINIT=YES is the default for APPN nodes. XCFINIT=YES is not allowed for pure subarea nodes.
- If XCFINIT=DEFINE is specified, VTAM joins the ISTXCF group at VTAM initialization. If the node is a pure subarea node, as it becomes aware of each other member node of the Sysplex, VTAM builds a TRLE definition for that node, but does not activate it. While this option does not enable SNA communications over XCF for pure subarea nodes, it does allow TCP/IP communications through XCF to be enabled. If the node supports APPN, as it becomes aware of each other member node of the sysplex, VTAM builds a TRLE and APPN PU definition for that node, but does not activate the APPN PU. This option is useful in APPN configurations where XCF connectivity is desirable for TCP/IP communications but not for APPN communications.
- If XCFINIT=NO is specified, VTAM does not join the ISTXCF group at VTAM initialization. The operator can cause VTAM to join the ISTXCF group at a later time by activating the ISTLSXCF major node (using the V ACT,ID=ISTLSXCF command).

Activation of the APPN PU for VTAM connectivity through XCF requires that HPR support in the node be specified as (or defaulted to) HPR=RTP. [If HPR=(RTP,ANR) is specified, VTAM will still treat XCF links as having RTP-level HPR support, not just ANR-level support.]

When VTAM builds a dynamic APPN PU definition to support XCF connectivity to another VTAM node in the sysplex, it adds dynamic PU definitions to the ISTLSXCF major node as each VTAM joins the XCF group. The first four characters of the names of the PUs defined in this manner default to ISTP. The next two characters are the &SYSCONE value of this VTAM. The last two characters are the &SYSCONE value of the partner VTAM.

When VTAM builds a dynamic TRLE entry to describe the connectivity characteristics used for XCF connections, it adds the entry to the ISTTRL major node. The first four characters of the names of the TRLEs defined in the TRL major node default to ISTT. The next two characters are the &SYSCONE value of this VTAM. The last two characters are the &SYSCONE value of the partner VTAM.

Requirement: The XCF major node (ISTLSXCF) and the PU representing the XCF connection to the VTAM on the other side of the XCF connection must each be active for SNA connectivity between the two nodes to be established.

Example: As an example, consider a sysplex with two nodes, Node A and Node B. Each can be started as a pure subarea node or as an APPN node. The APPN node can have XCFINIT=YES, XCFINIT=NO, or XCFINIT=DEFINE. The pure subarea node can have XCFINIT=DEFINE or XCFINIT=NO. Each can be started at the z/OS V1R7 level or above or with a prior level of z/OS code (prior levels cannot have XCFINIT=DEFINE). Table 45 on page 351 shows the XCF connectivity between the two nodes for each of these combinations. To make the table more readable, the following points apply:

- If either node (or both nodes) specify XCFINIT=NO, no XCF connectivity is established. The node with XCFINIT=NO will not join the ISTXCF group and other VTAM nodes in the group will not be aware of that node's presence in the Sysplex and will not build support for XCF connectivity to that node.
- Nodes that specify XCFINIT=NO when started can subsequently be allowed to participate in the Sysplex group by activating the ISTLSXCF major node. At that point, pure subarea nodes behave as if XCFINIT=DEFINE had been specified, and APPN nodes behave as if XCFINIT=YES had been specified.
- Pure subarea nodes at the pre-V1R7 level default to XCFINIT=NO. No other XCFINIT value can be specified.

Therefore, in Table 45 on page 351, each node can be one of the following types:

- A V1R7 or later pure subarea node specifying XCFINIT=DEFINE (or defaulted to DEFINE)
- A V1R7 or later APPN node specifying XCFINIT=DEFINE
- A V1R7 or later APPN node specifying XCFINIT=YES (or defaulted to YES)
- A pre-V1R7 APPN node specifying XCFINIT=YES (or defaulted to YES)

Table 45. VTAM to VTAM connection example					
		Node B			
		V1R7 pure subarea XCFINIT= DEFINE	V1R7 APPN XCFINIT= DEFINE	V1R7 APPN XCFINIT= YES	Pre-V1R7 APPN XCFINIT= YES
Node A	V1R7 pure subarea XCFINIT= DEFINE	Result 1	Result 1	Result 1	Result 2
	V1R7 APPN XCFINIT= DEFINE	Result 1	Result 6	Result 6	Result 5
	V1R7 APPN XCFINIT= YES	Result 1	Result 6	Result 3	Result 3
	Pre-V1R7 APPN XCFINIT= YES	Result 7	Result 4	Result 3	Result 3

Results:

1. Node A and Node B have XCF TRLEs defined, but no XCF APPN PUs are defined. SNA connectivity is not available. TCP/IP stacks on the two nodes can establish connectivity.
2. Node A and Node B have XCF TRLEs defined. Node B has the XCF APPN PU defined and activated. The PU remains in Pending Request Contacted (PREQC) state until manually deactivated by the operator. SNA connectivity is not available. TCP/IP stacks on the two nodes can establish connectivity.

3. Node A and Node B have XCF TRLEs and XCF APPN PUs defined and activated. Each PU activation will complete. SNA connectivity will be established. TCP/IP stacks on the two nodes can establish connectivity.
4. Node A and Node B have XCF TRLEs and XCF APPN PUs defined. Node A has activated the XCF PU. The PU remains in Pending Request Contacted (PREQC) state until the operator either manually deactivates the PU on Node A or manually activates the PU on Node B. TCP/IP stacks on the two nodes can establish connectivity. If the PU on Node B is activated, SNA connectivity is established.
5. Node A and Node B have XCF TRLEs and XCF APPN PUs defined. Node B has activated the XCF PU. The PU remains in Pending Request Contacted (PREQC) state until the operator either manually deactivates the PU on Node B or manually activates the PU on Node A. TCP/IP stacks on the two nodes can establish connectivity. If the PU on Node A is activated, SNA connectivity is established.
6. Node A and Node B have XCF TRLEs and XCF APPN PUs defined, but the PU is not activated by either node. TCP/IP stacks on the two nodes can establish connectivity. If the XCF PU is activated on either Node A or Node B, SNA connectivity is established.
7. Node A and Node B have XCF TRLEs defined. Node A has the XCF APPN PU defined and activated. The PU remains in Pending Request Contacted (PREQC) state until manually deactivated by the operator. SNA connectivity is not available. TCP/IP stacks on the two nodes can establish connectivity.

You can define a model for dynamic XCF local SNA PUs in a model major node to override the XCF default values for the operands of the PU definition statement (for example, CPCP or TGP). Coding a TRLE operand signifies that the model is for XCF connections. Do not specify CONNTYPE=LEN on the XCF model definition. Only one model definition is in effect at a time. The first model activated is the one used for the dynamic PUs. DYNTYPE=XCF can be coded in a model major node definition.

You can change the default values for the first four characters of the dynamic PU names added to ISTLSXCF by specifying up to four characters as the name of the model PU. The default values for the first four characters of the dynamic TRLE names added to the ISTTRL major node can be changed by specifying up to four characters on the TRLE operand of the model PU.

Generic resources

A generic resource is a method of allowing multiple application programs to be known by a common name. For additional description of generic resources, see [“Generic resources function” on page 324](#).

Generic resources requirements

Any VTAM application program running on a VTAM APPN node that is connected to an MVS coupling facility structure identified by the STRGR start option can be known by a generic resource name. All VTAM nodes having access to the common generic resource coupling facility structure, along with the structure itself, make up a generic resources configuration.

- A generic resources configuration must be part of a sysplex environment.
- All VTAMs in a generic resource configuration must be connected to an MVS coupling facility in which the generic resource structure can be allocated. VTAM connects to the generic resource structure when it is defined in the active CFRM policy.

The generic resource coupling facility structure stores the data needed by VTAM to implement the generic resources function. This data includes mappings that show:

- For each generic resource name, a list that identifies the generic resource members currently using that name. This is called a generic resource mapping.
- For each LU in session with one or more generic resource members, the application program network name and the generic resource name for each generic resource member. This is called a partner LU mapping.

In defining generic resource structure in a CFRM policy, you specify the amount of coupling facility storage which should be allocated on behalf of this structure. For information about determining the amount of storage necessary, see [“Determining the size of the coupling facility structure” on page 346](#).

- VTAM Version 4 Release 2 or later is required for all VTAMs in a generic resources configuration, and all VTAMs must be defined as network nodes or end nodes. If an application program is both a generic resource and uses multinode persistent sessions, all VTAMs in the sysplex must be Version 4 Release 4 or later.
- If you are using subplexing (that is you have specified the XCFGRPID start option), all resources that share a generic resource name must reside within the same subplex. That is, they must be running on VTAMs that have been started with the same value for the XCFGRPID start option. In addition, the generic resource coupling facility structure within each subplex must have a unique name. To ensure this, the structure name specified by the STRGR start option is suffixed with the 2-character XCFGRPID start option value. You must ensure that each of these subplex structure names is defined in the active CFRM policy. For example, if you have two subplexes within your sysplex, one with XCFGRPID=11 and one with XCFGRPID=02, and your STRGR option specifies the structure name ISTMYGR, you must ensure that the CFRM policy defines both an ISTMYGR11 and an ISTMYGR02 generic resource structure. You can display the fully suffixed generic resource structure name accessed on a VTAM node by issuing the D NET,ID=VTAM command.
- If instances of a generic resource exist on an end node, CP-CP sessions must exist between the end node and its network node server.
- If instances of a generic resource exist in the domain of multiple network nodes, CP-CP sessions should exist between each pair of network nodes to ensure correct session setup.
- The OLU host of a session initiation to a generic resource must support name translation. The TRANSLAT start option must include USERVAR.
- A network node server is required to provide generic resource resolution for resources on end nodes within the generic resource configuration. The recommended APPN generic resource configuration consists of a network node and served end nodes connected to the same generic resource structure with a backup network node connected to the same structure. But a backup network node server that is not connected to the same generic resource structure can be used. This requires that the end nodes allow broadcast searches, by coding ENBCAST=YES on the NETSRVR list. See the ENBCAST operand on the NETSRVR statement for the VTAMLST VBUILD TYPE=NETSRVR list in [z/OS Communications Server: SNA Resource Definition Reference](#).

If the network node server currently in use by the end nodes resides outside the sysplex, session setups to generic resources residing on these end nodes can fail if the NETSRVR list is not defined correctly. Symptoms of the failure would be similar to the following, where TSOXX is a generic resource name:

```

IST663I  CD DSRLST REQUEST TO   CDRMX  FAILED, SENSE=087D0001
IST664I  REAL  OLU=netid.luname  ALIAS  DLU=netid.TSOXX
IST889I  SID = EC7F2F218E97F0CF
IST1705I  SORDER = ADJSSCP FROM START OPTION
IST1705I  SSCPORD = PRIORITY FROM START OPTION
IST894I  ADJSSCPS TRIED,FAILURE SENSE,ADJSSCPS TRIED,FAIL SENSE
IST895I  ISTAPNCP          08400007   CDRMA          087D0001
IST895I  CDRMB            087D000A   CDRMX           087D000A
IST314I  END

```

- VTAM uses default settings to control generic resource resolution. To modify the default settings for all or specific generic resources you must define a generic resource preference table using GRPREF operands on the GRPREF statement for the VTAMLST VBUILD TYPE=GRPREFS statement. For more information see [z/OS Communications Server: SNA Resource Definition Reference](#). In addition, see [“Initiating sessions using the generic resource name”](#) on page 357.

Generic resource mapping

A map of the application programs that are members of each generic resource is kept in a generic resource structure. When an LU initiates a session using a generic resource name, VTAM establishes the session with one of the members that is mapped to that name.

Figure 99 on page 354 shows how the mapping identifies the application programs that are using the different generic resource names. In this example, APPLA, APPLB, and APPLC are represented to the network as JOHN. APPLD is using the generic resource name SEAN.

MVS Sysplex

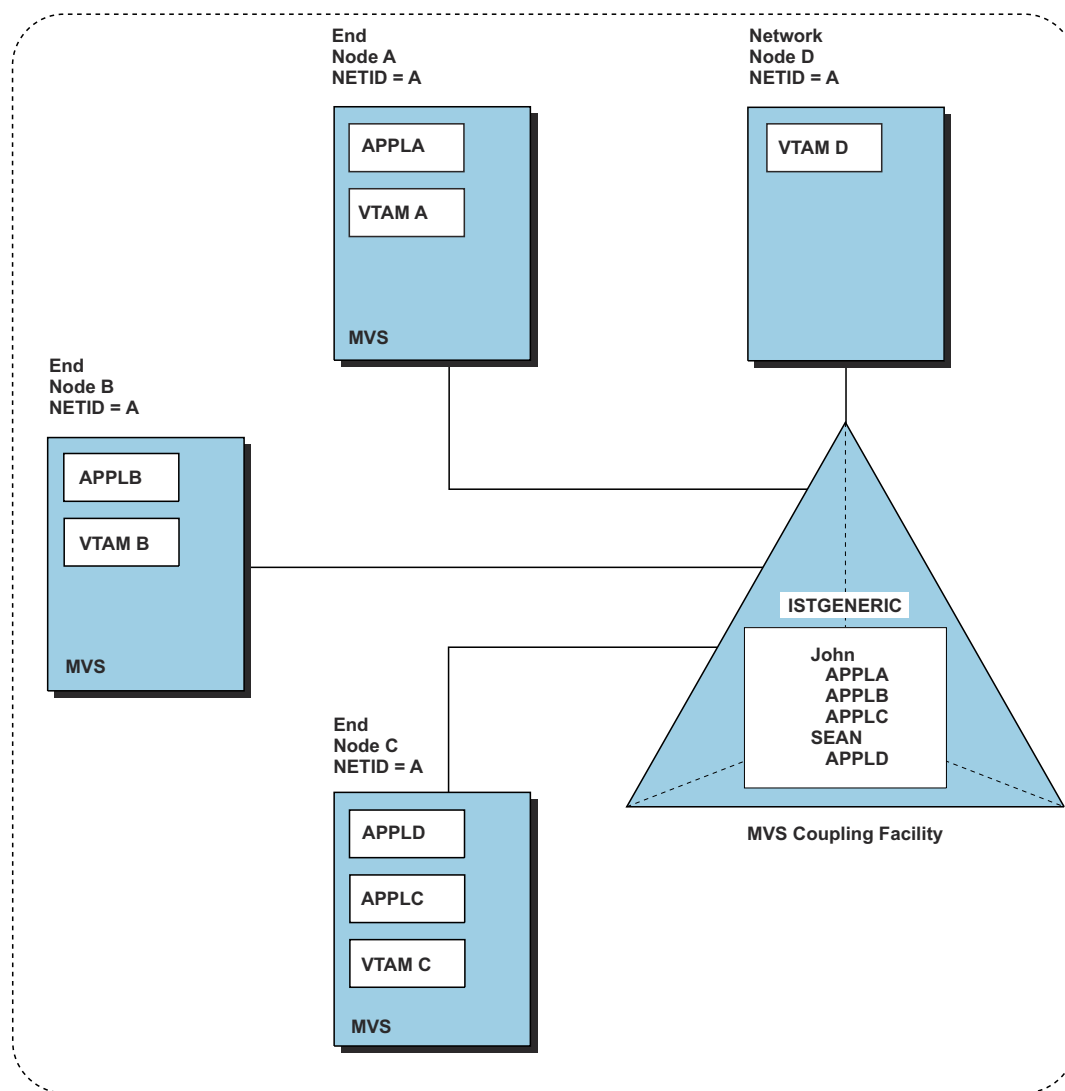


Figure 99. Generic resource mapping

Partner LU mapping

The generic resource structure also contains a mapping that identifies the partner LUs for each generic resource member. When the first session is established between an LU and a generic resource member, VTAM creates a mapping for that LU. This mapping, or affinity, identifies the name of the LU, the application program network name of the generic resource member, and the generic resource name. When subsequent parallel sessions are established using the same generic resource name, VTAM uses this mapping to ensure the session is established with the same generic resource member.

Figure 100 on page 355 shows partner LU mapping information. In this example, LU5 is in session with APPLB using the generic resource name JOHN and is in session with APPLD using the generic resource name SEAN. LU3 is in session with APPLA using the name JOHN.

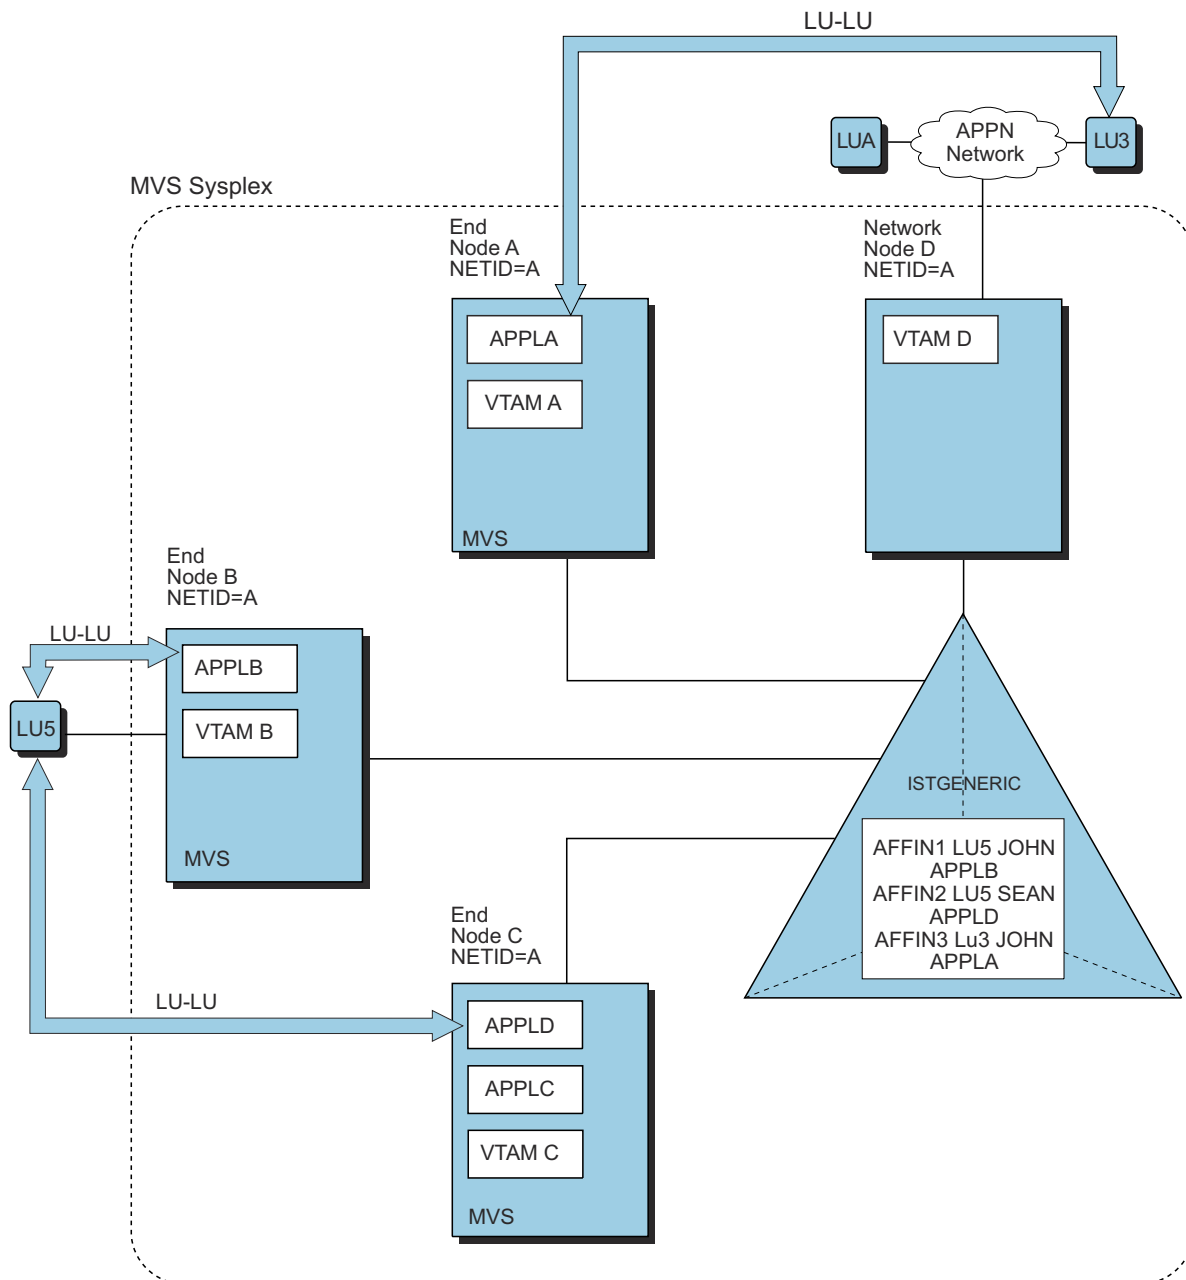


Figure 100. Partner LU mapping

When LUA initiates a session with JOHN, VTAM determines whether LUA has an affinity with a generic resource member for JOHN. Because LUA does not have an affinity with any generic resource member using the name JOHN, VTAM establishes the session with the generic resource member having the least number of LU-LU sessions (in this example, APPLC).

Figure 101 on page 356 shows the affinity that is created after LUA establishes a session with JOHN. VTAM resolves the session to APPLC; LUA is only aware of having established a session with JOHN. All subsequent parallel sessions between LUA and JOHN resolve to the same generic resource member, APPLC, regardless of session count.

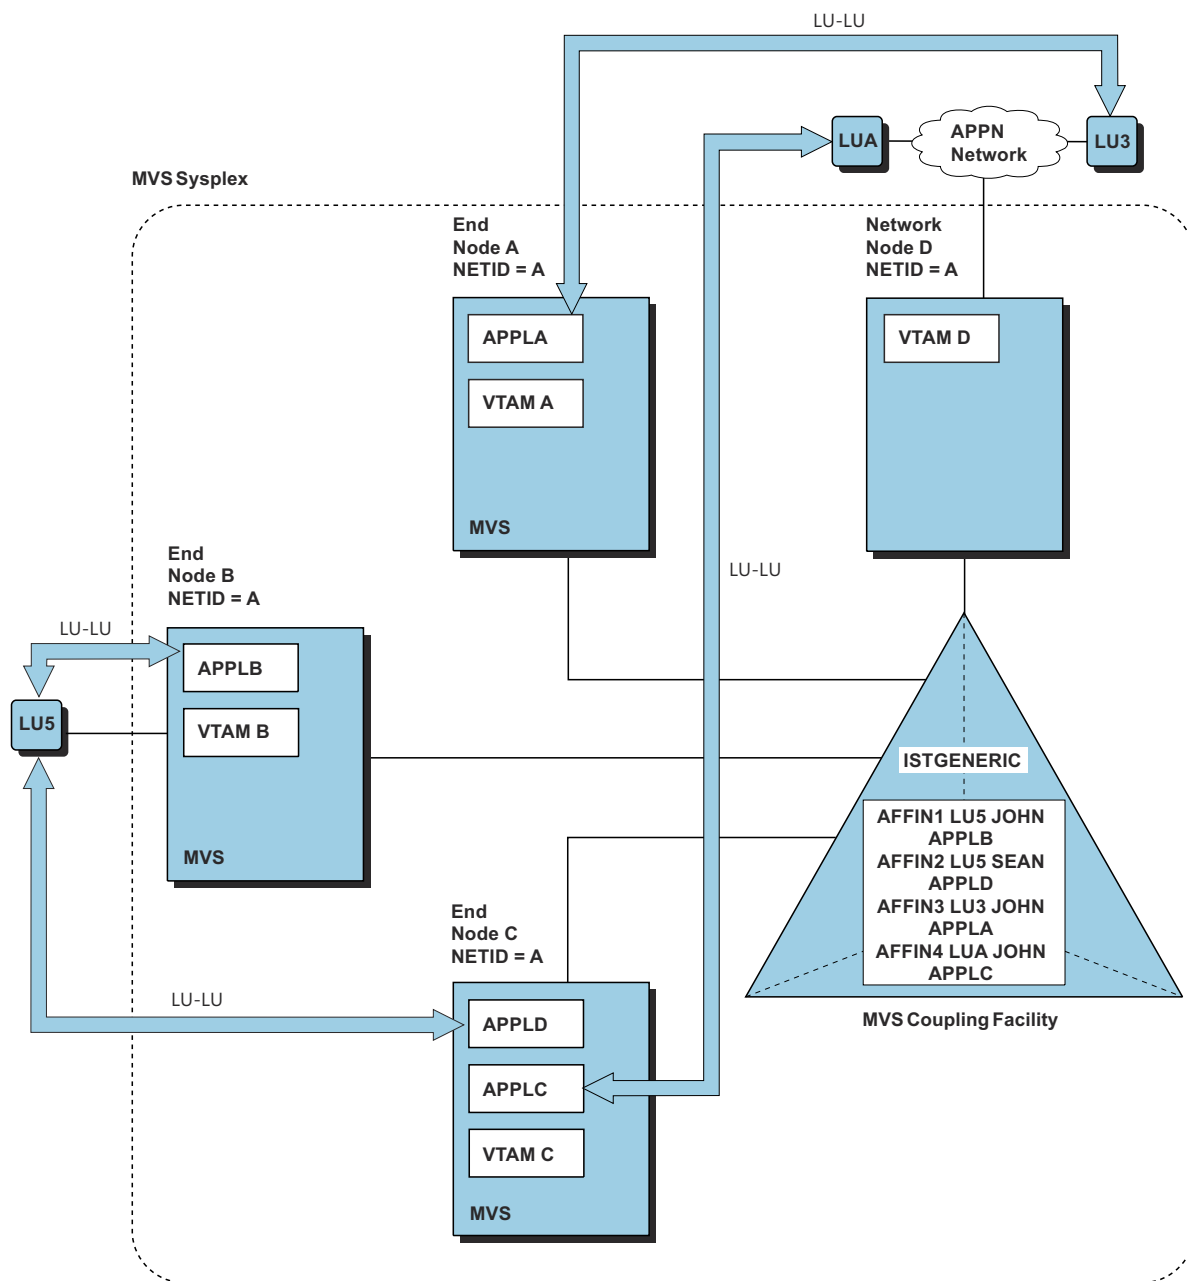


Figure 101. Session establishment with generic resource members

Affinities are deleted by their owner. Unless overridden by LUAFFIN=NOTAPPL specified for OPNDST, OPNSEC, or APPCCMD (CNOS, PREALLOC, or ALLOC) when a session is established, the application owns all affinities for sessions using:

- LU 6.2 synch point services
- LU 6.2 limited resource support
- LU 6.1 protocols

The application program is also owner of all affinities when AFFIN=APPL is specified for the SETLOGON OPTCD=GNAMEADD issued by the application, provided this is not overridden by LUAFFIN=NOTAPPL specified for OPNDST, OPNSEC, or APPCCMD (CNOS, PREALLOC, or ALLOC) when a session is established. The application program owns the affinity because of a particular session when LUAFFIN=APPL is specified for OPNDST, OPNSEC, or APPCCMD (CNOS, PREALLOC, or ALLOC) when a session is established.

The application program can delete an affinity by using the CHANGE OPTCD=ENDAFFIN macroinstruction, after all sessions between the application and the LU have terminated. The application program can

assign the affinity ownership to VTAM, even if there are still sessions with the LU, by issuing the `CHANGE OPTCD=ENDAFFNF` macroinstruction. VTAM will delete VTAM owned affinities when the last session between the LU and the application has terminated. LU 6.2 sessions, whose affinities are not owned by the application, are retained for 10 minutes after the last session, or for some other user-specified period (specified by start option `AFFDELAY`). For further information, see [z/OS Communications Server: SNA Programming](#).

Initiating sessions with generic resource members

Using generic resources, VTAM allows an application program to be known by a generic resource name. An application specifies its generic resource name on the `GNAME` operand of the `NIB` macroinstruction. The application establishes its association with the generic resource name on the `SETLOGON OPTCD=GNAMEADD` macroinstruction. VTAM uses RACF class `VTAMAPPL` to validate the generic resource name.

For any particular generic resource name, an LU can initiate a session using the generic resource name or the application program network name of the generic resource member, as shown in the following sections. For more information about how VTAM establishes sessions using the generic resources function, see [z/OS Communications Server: SNA Programming](#).

Initiating sessions using the application program network name

If an LU accesses a generic resource member by the application program network name, VTAM establishes the session with that member, regardless of the current workload of that member. In addition, when the first session between an LU and a generic resource member is started using the generic resource member application program name, any subsequent session initiation attempt from that LU to the generic resource name results in a session with one of the other generic resource members. For example, in [Figure 101 on page 356](#), if LUA initially used the application program name `APPLC` at logon, any subsequent session from LUA using the generic resource name `JOHN` is established with `APPLA` or `APPLB`.

Initiating sessions using the generic resource name

When an LU initiates a session using a generic resource name, VTAM performs generic resource resolution to identify a specific application name. VTAM determines whether the LU has an affinity, or mapping, with a generic resource member. If an affinity exists, VTAM establishes the session with the same generic resource member, regardless of current workload. (A session can be started with one of the other generic resource members using its application program network name.) If an affinity does not exist, VTAM uses the default or defined generic resource preferences from the generic resources preferences table to determine how to resolve the generic name. You can define a default generic resource preference or a generic resource preference for specific generic resources using the generic resources preference table (`VBUILD` type `GRPREFS`). See [z/OS Communications Server: SNA Resource Definition Reference](#) for the generic resource preference table.

The default generic resource resolution preferences for selecting a generic resource (unless overridden by defining a nameless entry in the `GRPREFS` table) are as follows:

- If a local application initiates a session to a generic resource, VTAM prefers generic resource members on the same VTAM node over those on other VTAM nodes. If no local generic resource members are active, then all generic resource members are eligible for resolution. This preference avoids the overhead associated with routing a session to other nodes. In many situations, you might want this local generic resource preference, but it can create a temporary or minor imbalance of session distribution. However, if the origin logical unit (OLU) is a session manager application or a `TN3270` server through which most of the sessions in your network are started, VTAM always selects a local instance of a generic resource; this can significantly overload a single instance of a generic resource. By setting the `LOCAPPL` generic resource preference in the `GRPREFS` table, you can control the preference for picking local instances of a generic resource.
- If a local LU that is part of a local SNA or local non-SNA major node initiates a session to a generic resource, VTAM prefers generic resource members on the same VTAM node over those on other VTAM

nodes. If no local generic resource members are active, then all generic resource members are eligible for resolution. This preference avoids the overhead associated with routing a session to other nodes. In many situations you might want this local generic resource preference, but it can create a temporary or minor imbalance in session distribution. However, if you have a large number of local LUs on the same host as the destination generic resource, VTAM always selects a local instance of a generic resource; this can significantly overload a single instance of a generic resource. By setting the LOCLU generic resource preference in the GRPREFS table, you can control the preference for picking local instances of a generic resource.

- If a generic resource resolution is from an LU that is not a local application or a local LU, then all generic resource members are eligible for resolution. However in the case of a third-party-initiated session (CLSDST-PASS) it might be preferable to select from generic resource members at the OLU CP for the session. If, for example, workload balancing had previously established the OLU of the session (for example, TN3270 using sysplex distributor), but the session setup path includes a CLSDST-PASS to a generic resource by a session manager on a non-OLU CP, you might prefer that the generic resource be resolved to a generic resource member on the OLU CP. This would avoid the overhead associated with routing the final session path to other nodes. By setting the PASSOLU generic resource preference in the GRPREFS table, you can control the preference for resolving to generic resource instances at the OLU CP of the session during CLSDST-PASS.
- VTAM calls the Workload Manager (WLM) to select an instance of the generic resource from all generic resource members that were previously determined to be eligible for this session. In most conditions, the workload at the time of the session initiation is the best way to determine how to resolve a generic name. However, you can set the WLM generic resource preference in the GRPREFS table to avoid calling the Workload Manager and instead resolve to a generic resource based on the fewest total number of active and pending sessions.
- IBM provides a default installation-wide generic resource exit routine (ISTEXCGR) that you can modify to select the member based on user-specified criteria. By default VTAM does not call the generic resource exit to select an instance of the generic resource. In most conditions, you can set generic resource preferences to customize generic resource resolution. However, you can cause the generic resource exit to be called using the GREXIT generic resource preference in the GRPREFS table.

The generic resource exit can use the information passed in the generic resource exit parameter list to identify an instance of the generic resource. This information includes:

- OLU name
- Generic resource name
- Eligible generic resource instances
- Session counts
- WLM's best generic resource instance
- Best session load balanced generic resource instance

The exit is called only if no affinity exists and after the other generic resource preferences have been used to identify eligible generic resource instances. Only the previously determined eligible generic resource instances are passed to the exit. For more information, see the generic resource resolution exit routine in [z/OS Communications Server: SNA Customization](#).

Implementation considerations

When implementing the generic resources function, the following considerations apply to the generic resources configuration:

- All generic resource members with the same generic resource name must reside within the same generic resource configuration. This means that all hosts where a generic resource member resides, and any resource selector nodes for that generic resource, must be in the same sysplex and be connected to the same coupling facility structure.
- LUs that initiate sessions using a generic resource name can be located outside the generic resource configuration that supports that generic resource name.

- Using two network node servers maintains generic resource resolution capability should one of the network nodes fail. A backup network node server outside the sysplex can be used. See the ENBCAST operand on the NETSRVR statement for the VTAMLST VBUILD TYPE=NETSRVR list in [z/OS Communications Server: SNA Resource Definition Reference](#).
- Performance can be enhanced by placing generic resource members at end nodes in the generic resource configuration. This frees up network nodes to provide routing services.
- An application program can be known by only one generic resource name.
- Generic resource members using the same generic resource name must have the same NETID.
- Generic resource names must be unique within a single network. A generic resource name cannot be identical to:
 - A USERVAR
 - An alias name
 - A real LU name
 - An ACBNAME
- The generic resources function supports multiple copies of the same application program; however, VTAM does not verify that each application program known by a generic resource name provides the same function.
- If your network has a security manager product installed, the security level for the generic resource name must be the same as the security level for all application programs that are using that generic resource name. The security level of the generic resource name is determined by the first generic resource member that associates itself with that generic resource name. If there is no security product, any application program residing on a VTAM that is connected to an MVS coupling facility can use any generic resource name.
- Using the generic resources function, VTAM distributes sessions across generic resource members rather than evenly across hosts. VTAM evaluates the current usage of each application program, regardless of the owning host. When activating generic resource members, consider the host capabilities. If VTAM A has twice the capability of VTAM B, VTAM A should support twice as many generic resource members.
- The MVS workload manager can also be used to balance allocation of sessions across hosts.

Coupling facility failures for generic resource configuration

There are three types of failures that can cause generic resource data to be lost:

- Failure of a VTAM node; all of the local data of the node is lost.
- Failure of a link between a VTAM node and generic resource structure; the VTAM node involved is unable to access data in the structure.
- Failure of generic resource structure; all of the data in the structure is lost.

Failure of a VTAM node

All local data in a VTAM node that needs to persist through a VTAM failure is replicated in the generic resource structure. If other VTAMs are present, they provide the necessary cleanup, and generic resources on the failing node can be restarted on another VTAM node. Affinities that should persist through a VTAM failure include LU 6.1 sessions, LU 6.2 sessions with SYNCPT synchronization support, multinode persistent sessions, and affinities claimed by the application through SETLOGON AFFIN=APPL, OPNDST/OPNSEC LUAFFIN, or APPCCMD LUAFFIN. If a failing VTAM has no local persistent data, a peer VTAM performing cleanup can delete the failed VTAM connection.

Note: If a VTAM connection to the generic resource structure is in failed-persistent state, an operator should not delete it.

Also, the generic resource structure remains allocated until all connections to it are deleted. Because there are no VTAMs to perform cleanup for the last VTAM to disconnect, you do need to delete the last connection to the generic resource structure to deallocate the structure.

Failure of a link between a VTAM node and generic resource structure

When a VTAM node loses its link to the generic resource structure, it initiates a structure rebuild to attempt to rebuild the generic resource structure on another coupling facility to which all VTAM nodes in the generic resource configuration have access. The rebuild structure process can be stopped by a peer node that is either unable to connect to, or loses its connection to, the new version of the generic resource structure.

If the generic resource structure is rebuilt to a coupling facility that is not accessible by all nodes, the VTAM that has no link to the current version of the generic resource structure disconnects from the structure. Other VTAMs provide the necessary cleanup, and when the failing link is repaired the disconnected VTAM attempts to reconnect to the generic resource structure.

Failure of the generic resource structure

When generic resource structure fails, each VTAM in a generic resource configuration is notified and attempts to initiate a structure rebuild to create a new version of the generic resource structure. The new version of the structure is replenished from the local data of each VTAM node in the generic resource configuration.

Removing a generic resource

After a generic resource name is defined, it can prevent other resources on VTAMs in the same sysplex from being defined with the same name. Therefore, there might be a need to delete a generic resource name in order to assign the name to another resource. One reason this might occur is because of the relationship between the USERVAR function and Generic Resources. USERVAR performs a similar function to generic resources and often customers migrate from USERVAR to generic resources. However, in preparation for migrating to use generic resources, a customer might choose to temporarily test a generic resource name before returning to using the identical USERVAR name. This would require the deletion of the generic resource name before the identical USERVAR name could be used.

When an application registers as a generic resource, several types of information related to the generic resource are created and maintained within VTAMs in the sysplex and within the generic resource coupling facility structure. This information includes a GENERIC USERVAR name, a generic mapping, and affinities.

- A GENERIC USERVAR name represents the generic name and is created on every VTAM in the sysplex that contains an application instance of the generic name, or has resolved the generic name.
- A generic mapping also represents the generic name, along with every application instance associated with the generic name. It is created in the generic resource coupling facility structure.
- Affinities are created when sessions are started to the generic resource. They relate specific LUs to specific application instances of the generic resource. There are two different types of affinities: application owned and VTAM owned. Application-owned affinities require action on the part of the generic resource application before they are terminated, while VTAM-owned affinities end when the corresponding session ends.

Deletion of a generic resource name requires the deletion of all of this information.

Steps for removing a generic resource

The generic resource includes a GENERIC USERVAR name, a generic mapping, and affinities. Only when all sessions have been terminated for all real instances, all affinities are deleted, and all real instances have performed CLOSE ACB processing, the generic resource can be successfully deleted. This following examples show how to delete a generic resource named neta.grappl, which has three real instances.

Before you begin

Procedure

Perform the following steps to remove a generic resource:

1. Identify all instances of the generic resource.

Issue `DISPLAY NET,ID=grappl,IDTYPE=GENERIC` at every host in the sysplex. The following sample output is generated:

```
d net,id=grappl,idtype=generic
IST097I DISPLAY ACCEPTED
IST075I NAME = GRAPPL, TYPE = GENERIC RESOURCE
IST1359I MEMBER NAME      OWNING CP   SELECTABLE  APPC
IST1360I NETA.NETAPPL1    SSCP2A      YES          NO
IST1360I NETA.APPL1       SSCP1A      NO           NO
IST1360I NETA.APPLAA1     SSCPAA      DEL          NO
IST1393I GENERIC RESOURCE NAME RESOLUTION EXIT IS ISTEXCGR
IST314I END
```

This display shows three real instances, each on a different VTAM, associated with the `grappl` generic resource.

- The `NETA.NETAPPL1` instance has issued an OPEN ACB and is selectable as a generic resource from any host in the sysplex.
 - The `NETA.APPL1` instance has issued an OPEN ACB but `SSCP1A` does not have CP-CP sessions with a network node server, and therefore, `NETA.APPL1` is not selectable from any host in the sysplex. However, it is still registered as a generic resource and is selectable for sessions that originate on `SSCP1A`.
 - The `NETA.APPLAA1` instance has either CLOSED its ACB, or remained OPEN and issued the VTAM API command to make itself not selectable.
2. Terminate all sessions with all instances of the generic resource.

Session information can be displayed using the VTAM commands `DISPLAY SESSIONS,LU1=real_instance_name,SCOPE=ALL` or `DISPLAY NET,ID=real_instance_name,SCOPE=ALL`.

Note: `DISPLAY NET,ID=real_instance_name` output indicates that the application is active (message `IST486I`), its generic name (message `IST1363I`), and its session partners (message `IST635I`).

```
d net,id=appl1,scope=all
IST097I DISPLAY ACCEPTED
IST075I NAME = NETA.APPL1, TYPE = APPL
IST486I STATUS= ACT/S, DESIRED STATE= ACTIV
IST1447I REGISTRATION TYPE = CDSERV
IST1363I GENERIC RESOURCE NAME GRAPPL REPRESENTS NETA.APPL1
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST1632I VPACING = 7
IST1938I APPC = NO
IST597I CAPABILITY-PLU ENABLED ,SLU ENABLED ,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = APPL1A
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = ECHO, STEPNAME = ECHO, DSPNAME = IST3C391
IST228I ENCRYPTION = OPTIONAL , TYPE = DES
IST1563I CKEYNAME = APPL1 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = 0 MAXIMUM = 0
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME      STATUS      SID      SEND RECV VR TP NETID
IST635I LAA3270A  ACTIV-S    EAABEE18F0E34F0F 0002 0001      NETA
IST314I END
```

```

d net,id=netappl1,scope=all
IST097I DISPLAY ACCEPTED
IST075I NAME = NETA.NETAPPL1, TYPE = APPL
IST486I STATUS= ACT/S, DESIRED STATE= ACTIV
IST1447I REGISTRATION TYPE = CDSERVR
IST1363I GENERIC RESOURCE NAME GRAPPL REPRESENTS NETA.NETAPPL1
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST1632I VPACING = 7
IST1938I APPC = NO
IST597I CAPABILITY-PLU ENABLED ,SLU ENABLED ,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = APPL2A
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = ECHO, STEPNAME = ECHO, DSPNAME = IST36CF1
IST228I ENCRYPTION = OPTIONAL , TYPE = DES
IST1563I CKEYNAME = NETAPPL1 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = 0 MAXIMUM = 0
IST171I ACTIVE SESSIONS = 0000000002, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME STATUS SID SEND RECV VR TP NETID
IST635I L2A3270A ACTIV-S F6ABEEC3A58FC5CF 0001 0000 NETA
IST635I L1A3270A ACTIV-S EAABEEC3A28FC3B2 0001 0000 NETA
IST314I END

d net,id=applaa1,scope=all
IST097I DISPLAY ACCEPTED
IST075I NAME = NETA.APPLAA1, TYPE = APPL
IST486I STATUS= ACT/S, DESIRED STATE= ACTIV
IST1447I REGISTRATION TYPE = CDSERVR
IST1363I GENERIC RESOURCE NAME GRAPPL REPRESENTS NETA.APPLAA1
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST1632I VPACING = 7
IST1938I APPC = NO
IST597I CAPABILITY-PLU ENABLED ,SLU ENABLED ,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = APPLAA
IST212I ACBNAME = FIRSTONE
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = ECHO, STEPNAME = ECHO, DSPNAME = ISTD5DC7
IST228I ENCRYPTION = OPTIONAL , TYPE = DES
IST1563I CKEYNAME = APPLAA1 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = 0 MAXIMUM = 0
IST171I ACTIVE SESSIONS = 0000000001, SESSION REQUESTS = 0000000000
IST206I SESSIONS:
IST634I NAME STATUS SID SEND RECV VR TP NETID
IST635I LAA3270B ACTIV-S EAABEE18F0E34F11 0001 0000 NETA
IST314I END

```

The DISPLAY NET,SESSIONS command can also show the session partners (message IST874I).

```

d net,sessions,lu1=netappl1,scope=all,list=all
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS
IST873I PLU SLU SID STATUS
IST874I NETA.NETAPPL1 NETA.L2A3270A F6ABEEC3A58FC5CF ACTIV
IST874I NETA.NETAPPL1 NETA.L1A3270A EAABEEC3A28FC3B2 ACTIV
IST924I -----
IST878I NUMBER OF PENDING SESSIONS = 0
IST924I -----
IST878I NUMBER OF ACTIVE SESSIONS = 2
IST924I -----
IST878I NUMBER OF QUEUED SESSIONS = 0
IST924I -----
IST878I NUMBER OF TOTAL SESSIONS = 2
IST314I END

```

3. Delete all affinities associated with all real instances of the generic resource.

Affinities associate a session partner to a generic resource and a specific member of the generic resource. They are created when the first session from a session partner is started to a generic resource. As long as the affinity exists, all subsequent sessions between the session partner and the generic resource will be directed to the specific member maintained in the affinity.

Affinities can be displayed using the `DISPLAY GRAFFIN,LU=*,*,GNAME=grappl` command.

```
d net,graffin,lu=*,*,gname=grappl
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = GENERIC AFFINITY
IST1706I PARTNER NAME      GENERIC RESOURCE    MEMBER    ATTRIBUTES
IST1707I NETA.LAA3270A      NETA.GRAPPL      APPL1      -VG--W--
IST1707I NETA.L1A3270A      NETA.GRAPPL      NETAPPL1    -VG--W--
IST1707I NETA.LAA3270B      NETA.GRAPPL      APPL1AA1    -AG6-W--
IST1454I 3 AFFINITIES DISPLAYED FOR LU=*,*
IST314I END
```

This display indicates the following about the affinities associated with `grappl`:

- The affinities for LAA3270A and L1A3270A are owned by VTAM (V).
- The affinity for LAA3270B is application owned (A) and persistent (6).
- All affinities indicate the associated sessions were started using the generic name (G).
- All affinities were created based on input from the MVS Workload manager (W).

All sessions between the session partner and the generic resource member must end before any associated affinities can be deleted. The different types of affinities would then be terminated using differing rules:

- An affinity that is VTAM owned is deleted when the last session between the session partner and the generic resource ends.
- An affinity that is application owned and persistent can be deleted only when the application notifies VTAM to delete the affinity. An affinity that is application owned, but not persistent, is deleted if the application closes its ACB or notifies VTAM to delete the affinity. Deleting persistent or application-owned affinities requires the application to be active and invocation of an application-unique command, or application termination. See application-specific documentation to determine how to delete application-owned and persistent affinities.

Sessions can be terminated by application users, application operator commands, or the VTAM command `VARY NET,TERM,LU1=real_instance_name`.

```
v net,term,lu1=laa3270a
IST097I VARY ACCEPTED
IST457I POSITIVE TERM COMMAND RESPONSE
IST455I LU1=NETA.LAA3270A SESSIONS ENDED
```

The following display shows that the VTAM-owned affinity for LAA3270A is deleted after the session is terminated.

```
d net,graffin,lu=*,*,gname=grappl
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = GENERIC AFFINITY
IST1706I PARTNER NAME      GENERIC RESOURCE    MEMBER    ATTRIBUTES
IST1707I NETA.L1A3270A      NETA.GRAPPL      NETAPPL1    -VG--W--
IST1707I NETA.LAA3270B      NETA.GRAPPL      APPL1AA1    -AG6-W--
IST1454I 2 AFFINITIES DISPLAYED FOR LU=*,*
IST314I END
```

4. Close the VTAM ACB for every real instance of the generic resource.

In some cases, VTAM retains information about the generic resource capability of the application as long as the application exists, so the application must perform CLOSE ACB processing. The

CLOSE ACB can be triggered by an application-unique command or the VTAM command VARY INACT, ID=*real_instance_name* can be used.

```
v net,inact,id=applaa1
IST097I VARY ACCEPTED
IST105I APPLAA1 NODE NOW INACTIVE
```

The command DISPLAY ID=*real_instance_name* can be used to verify that every instance of the generic resource has closed its ACB. This command needs to be issued at the owning host of each real instance to verify that the CLOSE ACB has been performed.

```
d net,id=applaa1
IST097I DISPLAY ACCEPTED
IST075I NAME = NETA.APPLAA1, TYPE = APPL
IST486I STATUS= INACT, DESIRED STATE= INACT
IST1447I REGISTRATION TYPE = CDSERV
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST1632I VPACING = 7
IST1938I APPC = NO
IST597I CAPABILITY-PLU INHIBITED,SLU INHIBITED,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = APPLAA
IST212I ACBNAME = FIRSTONE
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = ***NA***, STEPNAME = ***NA***, DSPNAME = ***NA***
IST228I ENCRYPTION = OPTIONAL, TYPE = DES
IST1563I CKEYNAME = APPLAA1 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = ***NA*** MAXIMUM = ***NA***
IST171I ACTIVE SESSIONS = 0000000000, SESSION REQUESTS = 0000000000
IST314I END
```

Note: Message IST486I indicates that the application is inactive.

5. Identify all hosts in the sysplex that contain a GENERIC USERVAR.

GENERIC USERVAR mappings can exist on a VTAM because a generic name was once resolved on this host. To identify whether a GENERIC USERVAR exists, issue DISPLAY RSCLIST, ID=grappl, IDTYPE=GENERIC on every host in the sysplex.

```
d net,rsclist,id=grappl,idtype=generic
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RSCLIST
IST1417I NETID      NAME      STATUS      TYPE      MAJNODE
IST1418I NETA      GRAPPL    ACTIV      GENERIC RESOURCE  **NA**
IST1418I NETA      GRAPPL    *NA*      GENERIC USERVAR   **NA**
IST314I END
```

This display indicates that the resource name GRAPPL with resource type GENERIC RESOURCE exists in the generic resource coupling facility structure. It also indicates that this VTAM has defined a type GENERIC USERVAR for GRAPPL.

6. Delete the generic resource information from the sysplex.

When all sessions have been terminated for all real instances, all affinities are deleted, and all real instances have performed CLOSE ACB processing, the generic resource *grappl* can be successfully deleted. Issue the command MODIFY GR, OPTION=DELETE, GNAME=*netid.grappl* at every host in the sysplex that has a GENERIC USERVAR for the generic resource. This should include every host where a real instance resided, but could also include other hosts in the sysplex.

All members of the generic resource are not selectable.

```
d net,id=grappl
IST097I DISPLAY ACCEPTED
IST075I NAME = GRAPPL, TYPE = GENERIC RESOURCE
IST1359I MEMBER NAME      OWNING CP      SELECTABLE  APPC
IST1360I NETA.APPL1       SSCP1A      DEL         NO
IST1360I NETA.NETAPPL1    SSCP2A      DEL         NO
IST1360I NETA.APPLAA1     SSCPAA      DEL         NO
```

```
IST1393I GENERIC RESOURCE NAME RESOLUTION EXIT IS ISTEXCGR
IST314I END
```

All affinities have been deleted.

```
d net,graffin,lu=*,gname=grappl
JOB 2 IST097I DISPLAY ACCEPTED
JOB 2 IST350I DISPLAY TYPE = GENERIC AFFINITY
IST1358I NO QUALIFYING MATCHES
IST1454I 0 AFFINITIES DISPLAYED
IST314I END
```

The command to delete the generic resource must be issued at every host in the sysplex that has a GENERIC USERVAR for the generic resource.

```
modify vtam,gr,gname=neta.grappl,option=delete
IST097I MODIFY ACCEPTED
IST223I MODIFY GR COMMAND COMPLETED
```

If a MODIFY GR,DELETE command is issued at a host that has temporarily lost connectivity to the generic resource coupling facility structure, then generic resource data local to the host can still be deleted. However the MODIFY GR DELETE command must complete successfully from at least one host connected to the generic resource coupling facility structure in order to delete the generic resource data in the structure.

When the generic resource name has been deleted, another resource such as an application or USERVAR can be defined with the same name. For instance, the MODIFY USERVAR command can be used to create a USERVAR for the name grappl.

```
modify vtam,uservar,id=grappl,option=update,value=appl1
IST097I MODIFY ACCEPTED
IST825I USERVAR DEFINED - NAME = GRAPPL, VALUE = APPL1
IST314I END

d net,id=grappl
IST097I DISPLAY ACCEPTED
IST075I NAME = APPL1, TYPE = APPL
IST113I GRAPPL IS A USERVAR WITH VALUE APPL1 IN NETWORK NETA
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1447I REGISTRATION TYPE = CDSERVR
IST599I REAL NAME = NETA.APPL1
IST977I MDLTAB=***NA*** ASLTAB=***NA***
IST861I MODETAB=***NA*** USSTAB=***NA*** LOGTAB=***NA***
IST934I DLOGMOD=***NA*** USS LANGTAB=***NA***
IST1632I VPACING = 7
IST1938I APPC = NO
IST597I CAPABILITY-PLU ENABLED ,SLU ENABLED ,SESSION LIMIT NONE
IST231I APPL MAJOR NODE = APPL1A
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST271I JOBNAME = ECHO, STEPNAME = ECHO, DSPNAME = ISTB972A
IST228I ENCRYPTION = OPTIONAL , TYPE = DES
IST1563I CKEYNAME = APPL1 CKEY = PRIMARY CERTIFY = NO
IST1552I MAC = NONE MACTYPE = NONE
IST1050I MAXIMUM COMPRESSION LEVEL - INPUT = 0, OUTPUT = 0
IST1633I ASRCVLM = 1000000
IST1634I DATA SPACE USAGE: CURRENT = 0 MAXIMUM = 0
IST171I ACTIVE SESSIONS = 0000000000, SESSION REQUESTS = 0000000000
IST314I END
```

Note: Message IST113I indicates that GRAPPL is a USERVAR for APPL1.

In general, you should use generic resources because they have a greater degree of availability and workload distribution. However, if you use the MODIFY GR,DELETE command to consolidate a group of generic resource applications to a single application represented by a USERVAR, then caution should be taken to ensure that the single USERVAR application can handle the consolidated workload.

Routine maintenance for VTAM nodes

Before bringing a VTAM down for routine maintenance, it is recommended that any existing affinities be ended. VTAM automatically ends most affinities when halted or during a failure. However, for LU 6.1 and

synch point-capable sessions, and application-owned affinities, VTAM maintains affinities between LUs and their generic session partners at both of the following locations:

- The node of the generic session partner
- The coupling facility

To end these affinities a CHANGE ENDAFFIN must be issued. See [z/OS Communications Server: SNA Programming](#) for information about the CHANGE ENDAFFIN macroinstruction. Not ending the affinities has the following consequences:

- A session attempting to use one of the affinities while the application's ACB is closed will fail.
- If the ISTGENERIC structure fails, the affinities will be lost.

Multinode persistent sessions

VTAM uses the coupling facility in the MVS sysplex to maintain session and HPR connection information for all multinode persistent session (MNPS) application programs. This allows VTAM to restore application sessions after instances of system or application failure, or as part of takeover of the application.

- For application program failures, an application can be restarted on the same VTAM on which it failed, through the single node persistent session (SNPS) support, or on a different VTAM through MNPS planned takeover or MNPS forced takeover.
- For planned or forced takeover or for operating system, hardware, or VTAM failures, a multinode persistent session application program can be restarted on any VTAM node that supports multinode persistent sessions in the sysplex, with the following considerations:
 - If a multinode persistent application program moves to the VTAM where one of its session partners is located, the sessions with that partner will not be restored. The sessions are terminated and must be restarted. This is also true if the multinode persistent application program moves to the VTAM that is the partner endpoint of the HPR connections associated with the application session.
 - If both session partners are located on the same VTAM when the failure occurs, the sessions between them cannot be restored. The sessions are terminated and must be restarted.
 - If the multinode persistent application moves to a VTAM node that is not connected to the multinode persistent session coupling facility structure, no sessions are restored.

Applications capable of using triple-DES encryption lose the triple-DES sessions if the application is moved to a node running a level of Communications Server for OS/390® that does not support triple-DES sessions. Also, if the application requires triple-DES encryption (ENCRTYPE = TDES24 was coded on the definition statement), the application is not allowed to successfully issue OPEN ACB on a node running a level of Communications Server for OS/390 that does not support triple-DES.

Until the application program is restarted, incoming data is maintained at the other end of the HPR connection. When the application is restarted and before the sessions are restored, VTAM stores the incoming data.

VTAM uses Control Vector 29 to preserve the session state. This includes the last sequence number, RH, and the first five bytes of the RU for each expedited and normal PIU sent outbound and inbound. At restart, VTAM passes that information to the application to resynchronize with the session partner. It is recommended that you check your application to verify that it makes use of this information.


The backup of a multinode persistence-enabled application program is shown in [Figure 102 on page 367](#).



1. APPL1 is in session with APPL2. The session is multinode persistent session (MNPS) capable and enabled. APPL1 has specified persistent on the APPL definition statement. Session data is being stored in the MNPS coupling facility structure.



2. When VTAM 1 fails, another VTAM in the sysplex detects the error. The application is marked pending recovery. Each VTAM connected to the MNPS structure starts a timer (PSTMER value). If the timer expires before recovery is successful, one of the VTAMs in the sysplex performs cleanup of the application program's session data in the MNPS structure.



3. Either through an operation or the automatic restart manager (ARM), APPL1 can be restarted on another multi-node persistent session capable VTAM.

In this case, APPL1 has been restarted on VTAM3.

The sessions are then restored using the data saved in the coupling facility structure.

Figure 102. Application program backup using multinode persistent sessions

Multinode persistent session support is enabled if single node persistence has been enabled, PERSIST=MULTI has been specified, and the configuration requirements are met. Using the PERSIST operand on the APPL definition statement, you can make an application program multinode persistence-capable (however, single node persistence is always available).

To change an application program from multinode persistence to single node persistence (and vice versa) you must first close and deactivate the application, modify the APPL definition statement (PERSIST operand), reactivate the application program major node, and start the application program.

When operating on a network node, use a model definition for your MNPS applications, because PERSIST=MULTI cannot be coded on a nonmodel application definition if VTAM is a network node.

MNPS forced takeover processing requires some additional settings at the taking over application to indicate that this is a forced takeover attempt. In particular, PARM=(FORCETKO=YES) must be coded on the ACB macroinstruction. Similarly, additional action is required at the application being taken over to indicate that forced takeovers are acceptable. In particular, PARM=(FORCETKO=ALL) or PARM=(FORCETKO=MULTI) must be coded on a SETLOGON OPTCD=PERSIST macroinstruction.

Multinode persistent session configuration requirements

Note: HALT and HALT QUICK override multinode persistent session capability of all application programs on the VTAM where the command is issued. HALT CANCEL will *not* override multinode persistent session capability. VARY INACT,TYPE=FORCE overrides multinode persistent session capability for a specific application program.

Following are the requirements for implementing multinode persistent sessions:

- A persistent application program must run on a VTAM node that is part of a sysplex environment.
- VTAM nodes in the sysplex wanting to support multinode persistent application programs must be running under a release of MVS that supports coupling facility services. The coupling facility must be using, at a minimum, coupling facility control code (CFCC) level 1.
- All VTAMs in the multinode persistent session (MNPS) configuration must be connected to an MVS coupling facility in which the MNPS structure can be allocated. The MNPS coupling facility structure stores the data and session information needed by VTAM to recover a persistent session.

For any MNPS using cryptography (and assuming that ICSF is active on all nodes where the MNPS application can operate), each z/OS Communications Server node that can recover an MNPS application should have the same host master key defined so a session key enciphered at one host is valid at the recovery host.

- In defining the structure in a CFRM policy, you specify the coupling facility storage that should be allocated on behalf of this structure. For information about determining the amount of storage necessary, see [“Determining the size of the coupling facility structure”](#) on page 346.
- All VTAM nodes providing multinode persistent session support must be defined as Rapid Transit Protocol (RTP) level nodes (HPR=RTP start option).
- z/OS Communications Server V1R6 or higher is required to implement MNPS forced takeover.
- If you are using subplexing (that is, you have specified the XCFGRPID start option), all MNPS application recoveries must be within the same subplex. That is, you must specify your ARM policy such that, when recovering a failed MNPS application, the new MNPS application instance is started within the same subplex (on a VTAM node that specified the same XCFGRPID start option value) as the original MNPS application. In addition, the MNPS coupling facility structure within each subplex must have a unique name. To ensure this, the structure name specified by the STRMNPS start option is suffixed with the 2-character XCFGRPID start option value. You must ensure that each of these subplex structure names is defined in the active CFRM policy. For example, if you have two subplexes within your sysplex, one with XCFGRPID=11 and one with XCFGRPID=02, and your STRMNPS option specifies the structure name ISTMYMNPS, you must ensure that the CFRM policy defines both an ISTMYMNPS11 and an ISTMYMNPS02 MNPS structure. You can display the fully suffixed MNPS structure name accessed on a VTAM node by issuing the D NET,ID=VTAM command.
- If using APPC/MVS as a persistent application, you must specify the PERSIST keyword. PERSIST has been added to the LUADD statement in the APPCPMxx parmlib member. When you code this keyword, APPC/MVS issues a persistent close when LUDEL is issued for one of its LUs.
- To receive the full benefit of multinode persistent session support, meshed (meaning every node has a one-hop connection to every other node in the sysplex) connectivity should exist between nodes in the sysplex. Also, all sysplex network nodes, even if no multinode persistent session applications plan to operate on the network node, should support RTP tower to guarantee recoverability of sessions.

In addition, if you plan to operate MNPS applications on the network nodes in the sysplex, you need to extend HPR into the network beyond the sysplex. In particular, all nodes physically adjacent to the network node or nodes where the application can operate should be configured to perform RTP tower function; otherwise, sessions which use paths through the non-RTP tower nodes will not be recoverable, because there will be no underlying HPR connection.

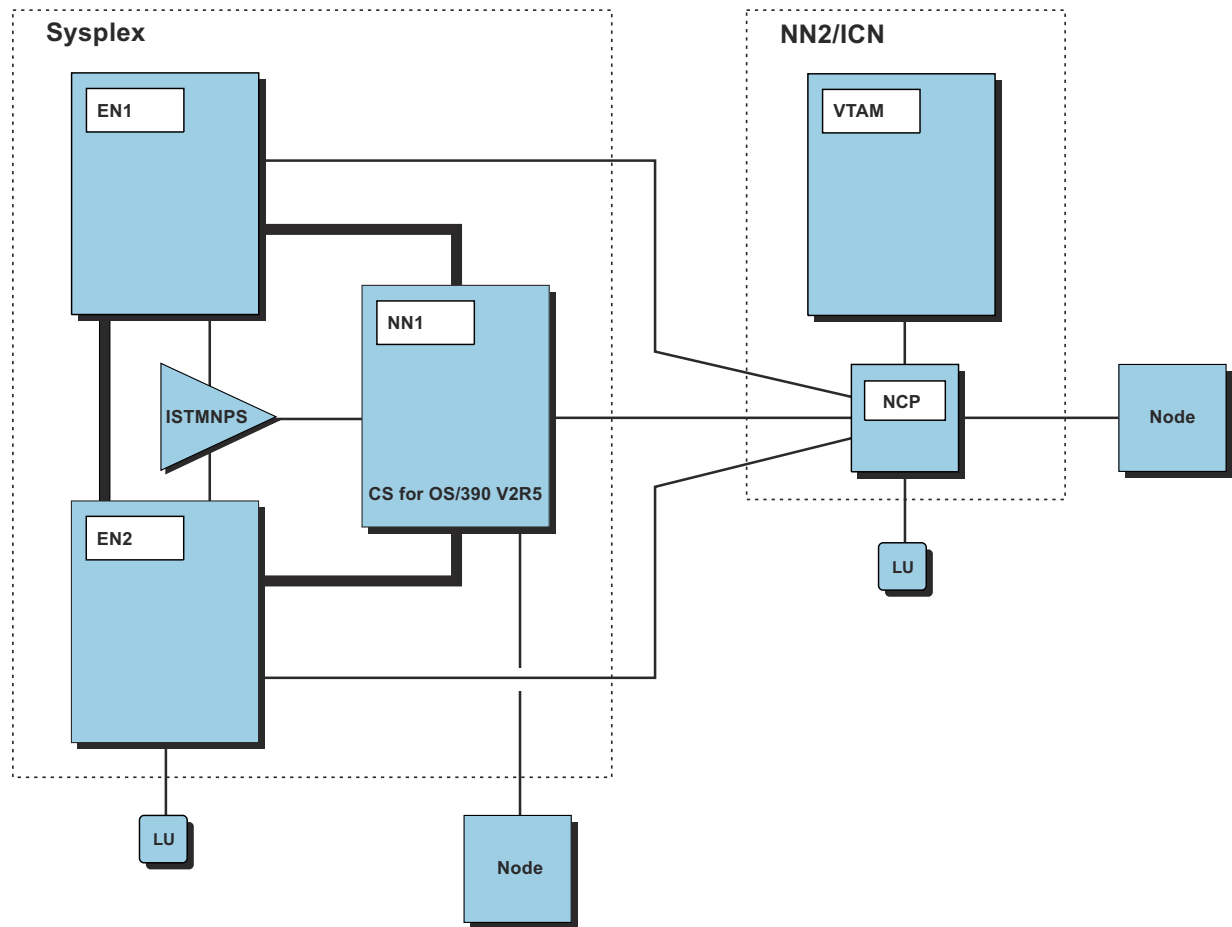


Figure 103. Multinode persistent session network example

Figure 103 on page 369 shows an example network configuration supporting multinode persistent sessions. EN1 and EN2 are both connected to the coupling facility. They are using the default name for the multinode persistent sessions structure, ISTMNPS. Also, in the sysplex is NN1, which is the network node server for the two end nodes. Outside of the sysplex is an interchange node, NN2, which has connections to all of the VTAMs in the sysplex. EN1, EN2, and NN1 are all RTP-capable nodes. It is optional for NN2 to be RTP capable; however, because NN1 is connected to the ISTMNPS structure and MNPS applications are anticipated to operate on NN1, NN2 should be configured to be RTP-capable as well. This will allow HPR connections to be established from NN1 to NN2 for sessions involving MNPS applications on the network node.

Note: If you plan to use VR-based TG connectivity in the sysplex, it is recommended that you also define at least one non-VR-based TG connection between any end node supporting MNPS application programs and its adjacent nodes. This will allow MNPS processing to calculate session paths to the end node that uses the non-VR-based TG and increase the probability that an HPR connection is used. See [“Session route setup at an end node”](#) on page 371 for more information.

Using multiple coupling facility structures for multinode persistent sessions

Multiple coupling facility structures can be defined to increase the number of multinode persistent sessions supported in the sysplex. Multiple structures can be defined on a single coupling facility or

across multiple coupling facilities. However, defining multiple structures on one coupling facility increases coupling facility overhead and is not advised.

To implement multiple multinode persistent session structures, define the structures in the MVS policy for each associated coupling facility. VTAM attempts to connect to a base structure and any alternate structures when they are defined in the CFRM policy. See [“Setting up the sysplex environment for VTAM and TCP/IP functions” on page 341](#) for information about defining structures used for multinode persistent sessions.

The structure name defined by the STRMNPS start option is the base structure. All other structures used for multinode persistent support are called alternate structures. The name of any alternate structure must begin with the name of the base structure followed by any number of unique and valid characters up to a total of 16 characters. For example, if you specify STRMNPS=ISTMNPS, an alternate structure name could be ISTMNPS01. Care should be taken when using subplexing to avoid generating an MNPS structure name for use in one subplex that might be mistaken as an alternate MNPS structure for a VTAM in another subplex. For example, if a VTAM is started in subplex 11 (that is, start option XCFGRPID=11 is specified), and STRMNPS=ISTMYMNPS is specified, and another VTAM in the sysplex is started with the default subplex (that is, start option XCFGRPID is not specified) and STRMNPS=ISTMYMNPS is specified, then the structure name generated for the MNPS structure in subplex 11 will be ISTMYMNPS11, which will appear to the second VTAM to be an alternate structure to its MNPS structure, ISTMYMNPS. See [z/OS Communications Server: SNA Resource Definition Reference](#) for additional information concerning the STRMNPS start option.

You can use the DISPLAY STATS command to view information about the base structure and its associated alternate structures.

Each structure on a coupling facility must have a unique name. If you are using multiple structures, the number of characters used to specify the name of the base structure (STRMNPS start option) must not be too large (to allow unique names for all the alternate structures).

Establishing multinode persistent sessions

When the application opens its ACB, it must indicate that persistent sessions are supported (PARMS=(PERSIST=YES)). The application must also include PERSIST=MULTI on its application definition. An MNPS application must be defined using a model application definition if the application operates at a network node. Upon completion of the OPEN ACB, the following information is added to the multinode persistent session structure:

- Name of the owning VTAM
- Application program in persistence disabled state
- Capabilities of the application (for example: NQN, APPC, and MACRF)

The application program then issues SETLOGON OPTCD=PERSIST to enable persistent LU-LU session support. If an application program wants to disable support for multinode persistent LU-LU sessions, SETLOGON OPTCD=NPERSIST should be issued. In both cases, the application program state in the coupling facility is updated to show "persistence enabled" or "persistence disabled," respectively.

You can issue the DISPLAY ID command to determine if application program sessions are recoverable.

- If the command is issued from the owning VTAM, the /M status modifier in message IST635I indicates that the session might be recoverable by multinode persistent session (MNPS) support. The current MNPS state is also displayed.

```
IST1550I MNPS STATE = ENABLED
:
IST635I APPLAA1 ACTIV/M-P EAABEE1819963497 0000 0000 0 0 NETA
```

- If the command is issued from another VTAM connected to the coupling facility structure in the sysplex message, IST1549I is issued. Message IST1549I contains MNPS STATE=ENABLED if the application program is recoverable by MNPS support.

```
IST1549I OWNER = NETA.SSCP1A MNPS STATE = ENABLED
```


Because sessions established for multinode persistent application programs use separate HPR connections from sessions not involving MNPS application programs, you might notice an increase in the number of HPR connections.

Session route setup at an end node

During session establishment, the end node provides a list of TG tail vectors to the network node server used in the computation of the session route. Because HPR is required for the recovery of a multinode persistent session, the EN tries to ensure the session route will transverse an HPR connection, by only providing to the NN server the endpoint TGs that connect the EN to RTP-capable nodes. This guarantees that the portion of the session path ending at the end node is HPR-capable and, therefore, the session is potentially recoverable. This might result in indirect routes when routing outside the sysplex. For example, in [Figure 104 on page 371](#), assume that APPL1 wishes to establish a session with APPL2.

- Without multinode persistent session enabled, the session path chosen is directly between MDH1 and VTAM1.
- With multinode persistent sessions enabled, the session path chosen is MDH1 to NN1 to VTAM1. So, the route for the session between APPL1 and APPL2 is a two-hop route.

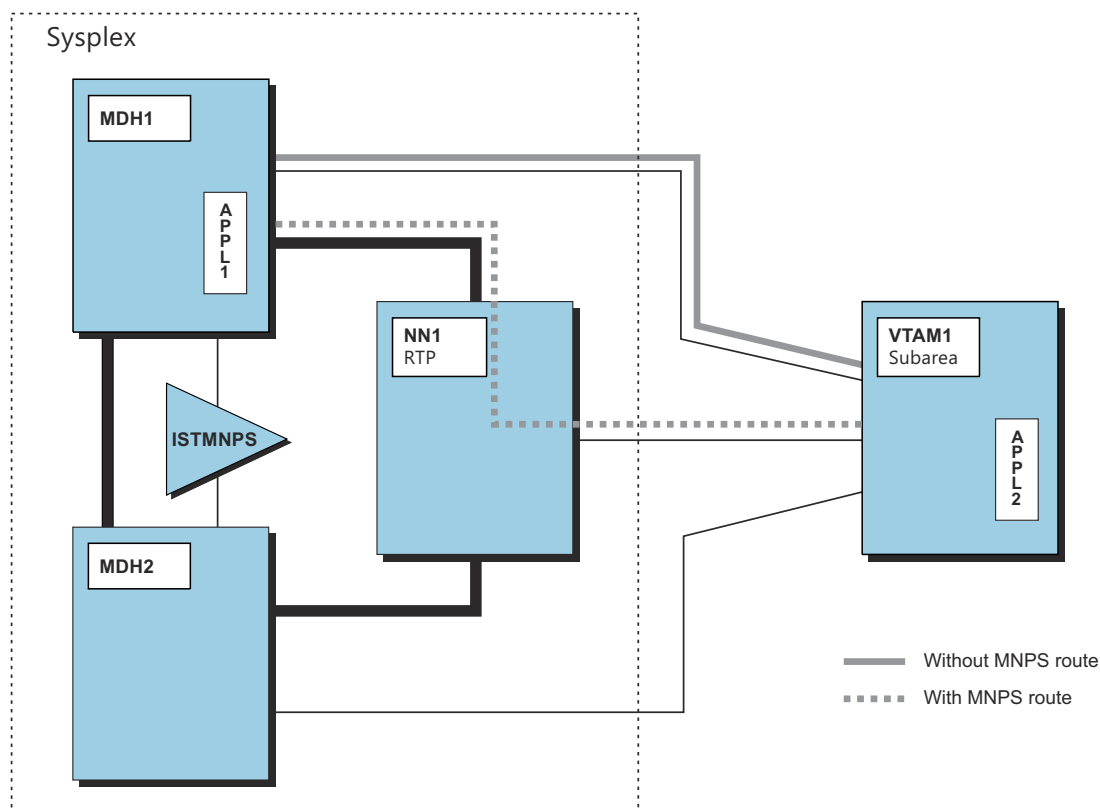


Figure 104. Session routes

Note: If APPL2 requests a session with APPL1 and VTAM1 sends the request directly to MDH1 using subarea protocols, the session is established without an HPR connection and is not recoverable. You might want to modify your search order to force the session requests for multinode persistent application programs to go through APPN over an HPR connection. See [“Controlling searches” on page 413](#) for additional information.

If during the APPN search no route is found, VTAM will establish the session using subarea routes if available. In this instance, the session is not recoverable even if the application program is multinode persistent-enabled.

After the session path has been established, VTAM can force a path switch to select a better session route. This is necessary because during the setup for the original path, the EN reports only HPR-capable links to the network node server. During session establishment:

- Any TG that connects an RTP-capable node to an ANR-capable node is omitted; however, these TGs might provide a better session route.
- VR-based TGs are not included in the subset of tail vectors if there are any other TGs that can be used. Recall that a VR-based TG can represent connections to either VTAM, NCP, or both. Also, note that at this time, the path to the partner LU is not known. If a VR-based TG is selected for the session path and if the partner LU is connected to the NCP in a composite network node, an HPR connection cannot be established because an NCP cannot be an RTP endpoint. Because the end node is trying to guarantee an HPR connection, the VRTG is not included because of the possibility of NCP being an endpoint.

For the better session route calculation, however, the VR-based TG can be included, because the other endpoint of the connection is now known and the network node server will only consider HPR-capable connections to that endpoint while trying to obtain a better session route.

A path switch is performed if an HPR-capable route is found and a different end node TG is selected for the route. After the path is determined and the session is established, VTAM places the appropriate information in the multinode persistent session coupling facility structure.

For example, in [Figure 105 on page 372](#) the session is established between APPL1 and APPL2 using EN1-NN1-NN3-EN3. The TGs considered when establishing the route were TGA, TGC, and TGD, because they guaranteed an RTP connection.

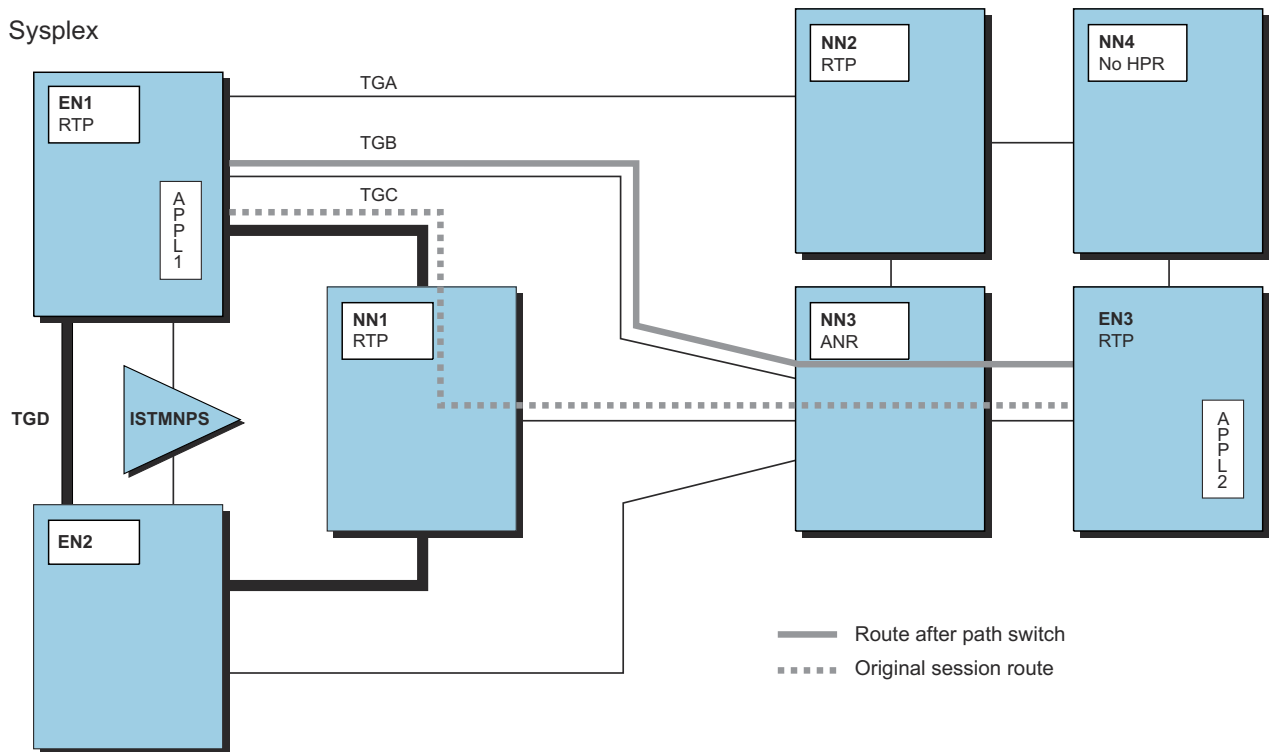


Figure 105. Path switch processing

After the session is established, EN1 attempts to find a better session route by initiating path switch processing. This time, TGB is included as a possible transmission group, and as a result the session route EN1-NN3-EN3 is selected. Because a different TG is used than in the original session route calculated, the data path for the HPR connection is switched to the better route.

If a better session path exists, and the data (or physical) path of the connection is switched to that path, the nonmultinode persistent partner is not told of the new path. The nonmultinode persistent partner remains unaware of the new path so that subsequent sessions will reuse the established HPR connection, because those sessions will, like the first, have their routes calculated using the subset of connections to RTP-capable nodes.

Because the nonmultinode persistent partner is not aware of the real data path, when the second session is set up, there appears to be an HPR connection already established for the route (EN1-NN1-NN3-EN3).

Therefore, the connection is reused. The MNPS EN, which is aware of the different physical path being used, will compare new session routes against the computed session path, and not the physical path, and will also be able to reuse existing HPR connections that traverse the same computed session path.

The `DISPLAY ID=RTP_PU_resource` command can be used to display the session path information known at a given node:

- A `DISPLAY ID=RTP_PU_resource` command, when issued at the multinode persistent node, provides both the computed session path and the physical path (if different) for the HPR connections used by multinode persistent application programs.
- A `DISPLAY ID=RTP_PU_resource` command provides the computed session path or actual path (being used) at the nonmultinode persistent session application program endpoint (the physical path).

Session route setup at a network node

Session establishment of multinode persistent sessions when the application is operating on a network node is less complicated than when the application is operating on an end node. The reason for this is that there is no concept of a list of TG tail vectors to report for a network node. All the transmission groups associated with a network node are known throughout the network, so any APPN network node in the network can calculate a session route to the MNPS NN using the network topology information. Because this information is known everywhere, there is no mechanism at the network node (such as TG tail vector lists) to force session paths to be calculated that will result in sessions that can be recovered by multinode persistent session processing.

To guarantee that session paths to the MNPS NN use HPR connections, it is required that the network configuration be managed in such a way that all nodes adjacent to the MNPS NN are RTP-tower capable. Additionally, all transmission groups connecting the MNPS NN to the adjacent RTP-tower capable nodes should be configured to support HPR connections. In such a configuration, regardless of the session path calculated, the portion of the path that terminates at the MNPS NN will consist of an HPR-capable connection to an RTP-tower capable node. This allows an HPR connection be established for at least that portion of the session path. The presence of the HPR connection allows the sessions using that connection to be recovered by multinode persistent session processing.

Coupling facility failures for multinode persistent session configuration

There are three types of failures which can cause multinode persistent data to be lost:

- Failure of a VTAM node; all local data for the node is lost.
- Failure of a link between a VTAM node and coupling facility structure; the VTAM node involved is unable to access or update data in the structure.
- Failure of a multinode persistent session (MNPS) structure; all of the data in the structure is lost.

Failure of a VTAM node

All other VTAMs connected to the MNPS structure will detect the VTAM failure. One of these VTAMs will process any data in the structure owned by the failing VTAM. For all MNPS application programs owned by the failing VTAM that are not enabled for recovery, all application data will be removed from the structure. For all MNPS application programs owned by the failing VTAM that are enabled for recovery, the application program will be marked recovery pending. This will allow the application program to recover its sessions on a different VTAM node. See [“Failure recovery processing” on page 374](#) for information about data recovery when a VTAM node fails.

Failure of a link between a VTAM node and multinode persistent structure

When a VTAM node loses its link to an MNPS structure and no system failure management (SFM) policy is in effect, it initiates a structure rebuild to attempt to rebuild the MNPS structure on another coupling facility to which all VTAM nodes have access. If an SFM policy is in effect, VTAM will be directed by MVS to either disconnect or start a rebuild. If the rebuild was started by VTAM because no SFM policy was in effect, the rebuild can be stopped by a peer node that is either unable to connect to, or loses its

connection to, the new version of the MNPS structure. A rebuild started by MVS because of SFM policy is never stopped by VTAM.

If the MNPS structure is rebuilt to a coupling facility that is not accessible by all nodes, any VTAM that has no connectivity to the current version of the MNPS structure disconnects from the structure. Other VTAMs will mark all applications owned by the disconnected VTAM as suspect. This indicates that the MNPS data is not up to date. When the failed link is repaired, the disconnected VTAM attempts to reconnect to the MNPS structure. If the connection is successful, the reconnecting VTAM will refresh all MNPS data for applications it owns from local data and reset the state to indicate the current persistent state of the application program as enabled or disabled.

Note: You should avoid performing planned takeover when a rebuild is underway, because the application will be viewed as suspect by the takeover VTAM, and the sessions will not be moved successfully.

Failure of an MNPS structure

When an MNPS structure fails, each VTAM connected to the structure is notified and attempts to initiate a structure rebuild to create a new version of the structure. The new version of the structure is replenished from the local data of each VTAM. Any data for applications that were recovery pending will be lost.

During a structure rebuild for an MNPS structure, data for the new version of the structure is always replenished from local data. In the case that an MNPS application program owning VTAM has failed and the state is recovery pending, if a rebuild occurs before the application program data is recovered, the data will be copied to the new version of the structure only if at least one VTAM has connectivity to both the original version of the structure and the new version of the structure. If this is not the case, the data is lost and the sessions cannot be recovered.

Failure recovery processing

When a VTAM connects to the multinode persistent session coupling facility structure in the sysplex, VTAM indicates that it wants to be notified when other connections to the structure end. The following describes the recovery processing for a persistent-enabled application program.

- First, the other VTAMs are notified that there is a node failure. One of the other VTAMs that is connected to the multinode persistent session structure marks the failing VTAM persistent-enabled applications as recovery pending.
- The other VTAMs then clean up the failing VTAM persistent-disabled application programs from the top of the coupling facility. When the coupling facility structure is cleaned up, no recovery can occur.

Note: From the perspective of the nonfailing partner, the sessions are still active. The sessions will remain active until the HPR connection is terminated.

- For all persistent-enabled application programs, a timer is started with the PSTIMER value specified by the application program. There is a timer for each MNPS application program residing in the failed VTAM. The PSTIMER indicates the amount of time an application can remain in recovery pending state. The recovering application must successfully issue an OPEN ACB before the timer expires on one of the VTAMs or VTAM cleans up the application program session information in the MNPS coupling facility structure.

Note: Affinities associated with generic resource application programs that are multinode persistent session-capable remain in the generic resource coupling facility structure until the timer expires. If the generic resource application program is not recovered, affinities are not deleted until the application program is restarted.

For recovery to occur, the application program must be persistent-enabled and the sessions must have traversed an HPR connection. If these conditions are met and there is a VTAM in the sysplex that is connected to the MNPS coupling facility structure, recovery can take place. The following describes the recovery process:

- An application is started either through an operator command or the automatic restart manager (ARM). The application opens an ACB with the same application program name as the application program

being recovered. Recovery can occur on the same VTAM that the application program resided on (if that VTAM has been restarted) or on a different VTAM.

- The recovering VTAM obtains the capabilities of the application being recovered from the coupling facility structure. The capabilities of the recovering application program must match those of the application program being recovered for recovery to continue.

Note: The use of model definitions to define the application program reduces the chances of capabilities not matching. The use of model definitions for MNPS applications is required at a network node.

See [z/OS Communications Server: SNA Programming](#) for details on the actual capabilities that must match.

- To indicate that the application has been recovered, the information in the coupling facility structure is updated to reflect the new owning VTAM and the application program is marked in persistence-ENABLED state.
- The new owning VTAM determines whether it has a CDRSC with the same name as the new application program. If a CDRSC definition exists and VTAM determines that the CDRSC represents the recovering application program when it was active on its previous owning VTAM, the sessions associated with the CDRSC are terminated and the CDRSC becomes a shadow resource. This must be done because an application program of the same name now exists.

If recovery occurs at a VTAM DLUS node, any sessions between the application and the DLUR-dependent LUs served by the VTAM DLUS network node are the exceptions to this processing. If the DLUS is intermediate or absent on the HPR connection path, the DLUS has only minimal session awareness of these sessions. These sessions are maintained and associated with the HPR connection that is to be rebuilt during MNPS recovery processing. Although the sessions are maintained, they are disconnected from the CDRSC representation of the application at this point in the processing and associated with the APPL RDTE that now represents the application. This allows OPEN ACB processing to continue.

- The new owning VTAM reads in all the application program data from the coupling facility into a VTAM data space.
- The new owning VTAM performs a path switch to reestablish a route for the session. During the path switch processing, the recovering side of the session will wait 90 seconds (or 4 minutes if both application programs are MNPS). If the connection is not reestablished, recovery is not successful.

The other end of the HPR connection is cleaned up when its path switch timer pops. For HPR connections being used by multinode persistent sessions, the minimum value for the path switch timer is four minutes. However, if a higher value is specified for the path switch timer (for VTAM, this is the HPRPST start option) at either session partner, it will be used.

Note: The values of the PSTIMER and HPRPST should be coordinated to synchronize effective use.

- The new owning VTAM re-creates the necessary session control blocks from the information in the coupling facility. The application program must restore the session, using the OPNDST RESTORE command, before session data traffic can resume completely. See the [z/OS Communications Server: SNA Programming](#) for additional information. Some applications might reset the sessions before resuming data traffic and if they do that, they are able to reduce the performance impact of MNPS by using NIBNTRCK on OPNSEC and OPNDST when establishing sessions.

Note: Applications using APPCCMD API will use the APPCCMD RESTORE command.

During recovery, you might notice a peak in storage usage on the new owning VTAM. When recovery is complete, storage usage should go back down.

A VARY INACT of an HPR connection being recovered because of multinode persistence processing will not be processed until recovery processing has completed.

Note: The OPEN will fail for an application program recovering for a failed generic resource application program if the recovering VTAM does not support generic resources. It will also fail if the recovery VTAM is attached to a generic resource structure name different from that of the MNPS application (the associated structure name is saved in the coupling facility).

MNPS planned and forced takeover processing

The two types of MNPS takeover processing that can be performed are:

- MNPS planned takeover is a more orderly takeover activity, forcing the target application to first enter a state, such as SNPS recovery pending, which indicates that recovery of some sort is necessary.
- MNPS forced takeover is a more dynamic takeover activity, in which the takeover application requests that the application to be taken over be shut down even in cases where the target application is currently active and operational.

During either type of takeover attempt, you might notice a peak in storage usage on the new owning VTAM. When takeover is complete, storage usage should go back down. A VARY INACT of an HPR connection being moved because of multinode persistence recovery processing will not be processed until takeover processing has completed.

Restriction: The OPEN will fail for an application program recovering for a failed generic resource application program if the recovering VTAM does not support generic resources or does not connect to the same generic resource structure.

MNPS planned takeover

For an MNPS planned takeover to occur, the application program must be in pending SNPS recovery state. An application enters this state when it has enabled persistence and then issues CLOSE ACB or fails. The state is maintained by the owning VTAM in the multinode persistent sessions coupling facility structure, so that other VTAMs in the sysplex are aware that the application is eligible for planned takeover processing.

Note: The D NET,ID command can be used to verify the state of any MNPS application within the sysplex.

If using APPC/MVS as a persistent application, you must specify the PERSIST keyword to support planned takeover. PERSIST has been added to the LUADD statement in the APPCPMxx parmlib member. When you code this keyword, APPC/MVS issues a persistent close when LUDEL is issued for one of its LUs.

The following describes the planned takeover process:

1. Either through an operator command or the automatic restart manager, an application opens an ACB with the same application program name as the application program that failed or closed its ACB.
2. The recovering VTAM obtains the capabilities of the application being recovered from the coupling facility structure. The capabilities of the recovering application program must match those of the application program being recovered for takeover to continue. See [z/OS Communications Server: SNA Programming](#) for details on the actual capabilities that must match.
3. The new owning VTAM determines if it has a CDRSC with the same name as the new application program. If a CDRSC definition exists and VTAM determines that the CDRSC represents the recovering application program when it was active on its previous owning VTAM, the sessions associated with the CDRSC are terminated and the CDRSC becomes a shadow resource. This must be done because an application program of the same name is attempting to move to this node.

If recovery occurs at a VTAM DLUS node, any sessions between the application and DLUR-dependent LUs served by the VTAM DLUS network node are the exceptions to this processing. If the DLUS is intermediate or absent on the HPR connection path, the DLUS has only minimal session awareness of these sessions. These sessions will be maintained and associated with the HPR connection that is to be rebuilt during MNPS recovery processing. Although the sessions are maintained, they will be disconnected from the CDRSC representation of the application at this point in the processing and associated with the APPL RDTE that now represents the application. This allows OPEN ACB processing to continue.

4. OPEN ACB processing is performed for the application at the takeover VTAM, but no response is posted to the application. VTAM must first obtain ownership of the application from the current owning VTAM before the OPEN ACB can be considered to be a success.
5. The VTAM attempting to acquire ownership sends a takeover request to the current VTAM (by XCF). If accepted, the current owning VTAM updates the multinode persistent session structure to show the takeover VTAM as the application owner, and then sends a positive takeover reply.

6. When the takeover VTAM receives the positive takeover reply, VTAM will claim the structure storage lists associated with the application and will mark the application as being enabled for persistence. The response to the OPEN ACB is then posted to the application, because this VTAM has now acquired ownership of the application.
7. After transferring ownership, the previous owning VTAM cleans up its local representations of the sessions, using care to prevent any modifications to the coupling facility data. The HPR connections associated with the sessions are also cleaned up locally, although the other endpoint is unaware of the processing as VTAM will not send, or handle any received, information about these connections. The image of the application is also cleaned up at the previous owning VTAM. Any sessions not involving HPR pipes are terminated at this point by the previous owning VTAM.

If the previous owning VTAM is a DLUS network node, minimal awareness of sessions between the application and DLUR-dependent LUs served by this VTAM DLUS will be maintained after the local session representation is cleaned up. This minimal session awareness must be kept to provide accurate session limit management for the dependent LUs. To maintain this session awareness, a CDRSC representation of the application must either be reused from the shadow set of resources (if a predefined CDRSC had existed for the application before OPEN ACB having been performed), or a CDRSC must be created dynamically. If no CDRSC can be reused or built (for example, CDRDYN=NO was specified as a start option), then session awareness cannot be maintained, and the session is terminated.

8. The takeover VTAM performs path switches to make this VTAM the new endpoint for the HPR connections associated with the application. During the path switch processing, the recovering side of the session will wait 90 seconds for the other endpoint to restart the HPR connection. If no response is received, recovery does not complete.

The other end of the HPR connection is cleaned up if its path switch timer pops. For HPR connections being used by multinode persistent sessions, the minimum value for the path switch timer (for VTAM, this is the HPRPST start option) is four minutes. However, if a higher value is specified on the HPRPST start option, it will be used.

9. The new owning VTAM re-creates the necessary session control blocks from the information in the coupling facility. The application program must restore the session, using the OPNDST RESTORE command or APPCCMD RESTORE (if using the APPCCMD API), before session data traffic can resume completely.

Note: The OPEN will fail for an application program recovering for a failed generic resource application program if the recovering VTAM does not support generic resources and connect to the same generic resource structure.

MNPS forced takeover

MNPS forced takeover processing can be performed when the application program is in a wider range of states, the common state characteristic being that the application is active and has enabled persistence. The application state is maintained by the owning VTAM in the multinode persistent sessions coupling facility structure. Other VTAMs in the Sysplex currently use the state information when processing OPEN ACBs or VTAM failure conditions, and now will use it to determine if the application is eligible for forced takeover processing.

In addition to state requirements, the application being taken over must indicate that it will accept forced takeover requests. This capability is expressed by coding PARMS=(FORCETKO=ALL) or PARMS=(FORCETKO=MULTI) on the SETLOGON OPTCD=PERSIST macroinstruction used to enable persistence in the first place. This capability is saved into the MNPS structure, much like the state of the application.

The D NET,ID command can be used to verify the current takeover capability of the application, both at the owning VTAM and at remote VTAMs connected to the same coupling facility structure. Any of the following messages can be generated as part of the D NET,ID command:

```
IST2061I NO FORCED TAKEOVER REQUESTS ARE ACCEPTABLE
IST2062I type FORCED TAKEOVER REQUESTS ARE ACCEPTABLE (where type can be MNPS or SNPS)
IST2063I ALL FORCED TAKEOVER REQUESTS ARE ACCEPTABLE
```


Requirement: Both the forced takeover and owning VTAMs must be z/OS Communications Server V1R6 level or higher.

The mechanics of an MNPS forced takeover are very similar to an MNPS planned takeover. The highlights of the forced takeover processing are:

1. An application opens an ACB with the same application program name as the target application program. The ACB macroinstruction must specify `PARMS=(FORCETKO=YES,PERSIST=YES)` in order for the OPEN ACB processing to treat this request as one that is permitted to initiate an MNPS forced takeover sequence. The application being taken over must also have indicated that MNPS forced takeovers are permitted, using the SETLOGON macroinstruction.
2. The taking over VTAM obtains the capabilities of the application being taken over from the coupling facility structure. The capabilities of the taking over application program must match those of the application program being taken over for forced takeover to continue. See [z/OS Communications Server: SNA Programming](#) for details on the actual capabilities that must match. In addition, for forced takeover, the state of the application being taken over must be one in which forced takeover is permitted, and the application being taken over must indicate that it will accept forced takeover requests.
3. The same CDRSC and DLUR processing is performed for forced takeover as is performed for planned takeover.
4. OPEN ACB processing is performed for the application at the taking over VTAM, but no response is posted to the application. VTAM must first obtain ownership of the application from the current owning VTAM before the OPEN ACB can be considered successful.
5. The VTAM attempting to acquire ownership sends a takeover request to the current VTAM (by XCF). This signal will carry a new indicator to say this is a forced takeover request. The target node, if the application status is unchanged or if an earlier forced takeover attempt from a different node has not already been accepted, will initiate CLOSE ACB processing for the application. This is the same persistent CLOSE ACB processing that was required before MNPS planned takeover request processing, but now, for forced takeover, the CLOSE ACB is driven internally by VTAM. This persistent CLOSE processing is still required, even for forced takeover, in order to get a stable representation of the application sessions and underlying HPR pipes for takeover purposes.

The application is notified of the persistent CLOSE processing using the TPEND exit. When the persistent CLOSE ACB completes, the current owning VTAM, just as in planned takeover, updates the multinode persistent session structure to show the taking over VTAM as the application owner, and then sends a positive takeover reply.

Note: Because the CLOSE ACB is now performed at the current owning VTAM as part of forced takeover processing, the user should anticipate a slightly longer recovery time for forced takeover as compared to planned takeover.

Recommendation: An application willing to accept MNPS forced takeover requests must consider carefully the PSTIMER setting chosen (if one is used at all). A persistent CLOSE ACB for the application is driven by VTAM as part of the forced takeover logic and after that has completed successfully, a non-persistent CLOSE ACB is driven. The application PSTIMER value should be set to a value at least long enough to allow the internal VTAM processing to occur between the completion of the persistent CLOSE and the start of the non-persistent CLOSE. An interval of 30 seconds or more is recommended for this situation.

6. When the taking over VTAM receives the positive takeover reply, VTAM will claim the structure storage lists associated with the application and will mark the application as being enabled for persistence. The response to the OPEN ACB is then posted to the application, because this VTAM has now acquired ownership of the application. Takeover processing proceeds along the same lines as planned takeover at this point.
7. After transferring ownership, the previous owning VTAM cleans up its local representations of the sessions and underlying HPR pipes in the exact same manner as it did for planned takeover processing. The local copy of the application is discarded when the sessions and pipes are eliminated.
8. As with planned takeover, when the new owning VTAM has re-created the necessary session control blocks from the information in the coupling facility, the application program must restore the session,

using the OPNDST RESTORE command or APPCCMD RESTORE (if using the APPCCMD API), before session data traffic can resume completely.

What to do if recovery does not occur or complete

The D NET,ID command can be used to display the sessions associated with an MNPS application. As described earlier, the session that is eligible for MNPS recovery is denoted with a /M status modifier; however, the presence of the /M status modifier does not guarantee that the session will recover if the application were to move or to fail. Factors, such as whether the application is enabled for persistence and whether the session partner is a DLUR-related dependent LU, for example, can affect the actual recovery of the session.

If the display information indicates that a session is not recoverable, and if a message is displayed indicating that a session is not recoverable, and you expected it to be recoverable, do the following steps:

- Verify that PERSIST=MULTI is specified on the APPL definition statement. If model application definition statements are being used to define your application program, verify that the application program is active and PERSIST=MULTI is coded on the application model.
- Verify that PARMS=(PERSIST=YES) is specified on the OPEN ACB for the application program.
- Verify that the coupling facility structure name matches the STRMNPS start option for all VTAMs.
- Verify that HPR=RTP start option is coded and not overridden by the definitions for the PU or TG.
- For end nodes, verify that at least one adjacent node has HPR=RTP start option specified.
- For network nodes, verify that the path chosen for the session is one that traverses an adjacent node with HPR=RTP coded (or the configuration option that applies to the adjacent node if not a VTAM).
- Verify that the adjacent CDRM (for migration data host node) does not favor subarea routing over APPN (SSCPDYN start option or SSCORD and SORDER).

Note: SSCORD and SORDER can be specified as start options or specified separately for each ADJSSCP table.

If you have a multinode persistent session-capable application program, and the sessions are not being recovered when you expect them to be recovered, do the following steps:

- Verify that the recovering application program is using the same name as the failed application program.
- Verify that an MVS system symbol is used as part of the value for LUAPFX. If model application definition statements are being used to define your MNPS application program and LUAPFX is specified, use an MVS system symbol as part of the value for LUAPFX to reduce collision of LUALIAS values.
- Verify that the sessions do not involve LUs active where the recovery is taking place or LUs active at the node where the MNPS application was previously active.
- Verify that the PSTIMER has not expired or that it is not too short. Recall that when the PSTIMER expires, the session information is cleaned up. If the timer value is too short, there might not be sufficient time to recover the sessions.
- Verify that the partner node still has sessions showing active. It could be that the MNPS node shows the sessions as PRECOVERY, but the HPR connection has timed out and sessions are cleaned up at the partner node.
- Verify that the sessions are not in PRECOVERY state. If they are, the recovery might be occurring but is slow to complete.
- Verify that the application program being recovered has PERSIST=MULTI specified on its APPL definition statement and PARMS=(PERSIST=YES) on its ACB statement.
- Verify that the new VTAM (where the application is being recovered) has HPR connectivity and (if VTAM is an end node) has CP-CP sessions to an NNS (which might be needed to calculate the new route to the partner endpoint). Use the DISPLAY NET,STATS,TYPE=CFS command to verify that the new structure has connectivity to the multinode persistent session coupling facility structure.
- Verify that if the attempt is a forced takeover, that the takeover and the current owning VTAM are at the correct level, which is at a minimum z/OS Communications Server V1R6. Verify

that the ACB macroinstruction of the application issuing the forced takeover request correctly specified `PARMS=(FORCETKO=YES)`. Verify that the application being taken over has correctly indicated the ability to accept forced takeover requests by specifying `PARMS=(FORCETKO=ALL)` or `PARMS=(FORCETKO=MULTI)` on `SETLOGON OPTCD=PERSIST` (and that no subsequent `SETLOGON OPTCD=PERSIST` specified `PARMS=(FORCETKO=NONE)` or `PARMS=(FORCETKO=SINGLE)`). Finally, if the application being taken over was once owned by a pre-z/OS Communications Server V1R6 node, then ensure that `SETLOGON OPTCD=PERSIST` with `PARMS=(FORCETKO=ALL)` or `PARMS=(FORCETKO=MULTI)` was performed after the application moved to a z/OS Communication Server V1R6 node.

- Verify that a CDRSC representation of the MNPS application could be created, or reused, at the VTAM DLUS node when moving the MNPS application to another VTAM node in the sysplex.

TSO generic resources

When TSO is part of the sysplex environment and it exists on multiple sysplex members, the use of TSO through VTAM can be extended to allow a generic name to be assigned to all TSO/VTAM application programs. A TSO/VTAM on one MVS system can be known by the same generic name as a TSO/VTAM on any other MVS system. All TSO/VTAM application programs sharing a particular generic name can be concurrently active.

Restriction: Use this configuration only if all the systems share the same security database. Otherwise, the user might not be able to reconnect to their session by using the generic resource, or know which password to enter when they log on.

Using a generic name allows for increased availability and workload balancing because all TSO/VTAM application programs can be accessed by the same generic name. For example, the MVS workload manager can be used to make efficient use of system resources by selecting a session partner based on balanced load distribution from among TSO/VTAM application programs with the same generic name.

An enhanced TSO reconnect function is also available; this function is dependent on a Job Entry Subsystem (JES) complex. A JES complex can be either a JES2 Multiaccess spool (JES2 MAS) or a JES3 processor complex. Either of these is a collection of host systems that share spool and other system facilities. For example, one piece of shared information is the identity and location of every TSO user address space. In a TSO/GR configuration, JES is asked for the system location of a disconnected TSO user when a reconnect logon request is submitted for the TSO generic name. The scope of visibility for JES is the JES complex. In this way, reconnect works in a predictable way for TSO/GR only when all members of the GR set are contained within one JES complex. This does not mean that you cannot have two, distinct TSO/GR sets with two, distinct JES complexes, but that arrangement is not recommended. Management is easier if you have only one JES complex for a parallel sysplex. Also, the TSO/GR set does not, and cannot, extend beyond the boundaries of a parallel sysplex.

When a generic name is used for TSO, unique names must be given to each of the TCAS, terminal control address space, APPL definition statements. The name *TSO* should not be used as the name on the TCAS APPL definition statement in any of the sysplex members. *TSO* is the TCAS ACBNAME in each sysplex member. Using *TSO* as the name on a TCAS APPL definition statement and as a generic resource member name can result in a local TCAS being selected rather than the generic member, whenever generic name resolution selects the member name *TSO*. The result is an unexpected distribution of TSO sessions in the generic resource environment.

To assign a generic name to a TSO/VTAM application, specify the `GNAME` parameter. The `GNAME` parameter can be:

- Coded in the `TSOKEYxx` parmlib member
- Specified on the `START` command

Sysplex-wide security associations

TCP/IP uses the coupling facility in the MVS sysplex to provide the ability to recover IPsec security associations for a distributed DVIPA on another node in the sysplex, either when a TCP/IP stack fails or

when a planned takeover is initiated. The configuration and use of this function is described in the [z/OS Communications Server: IP Configuration Guide](#).

Coupling facility failures for sysplex-wide security associations

There are four types of events that can impact the sysplex wide security associations function:

- [“Failure of a TCP/IP stack” on page 381](#)
- [“Failure of a VTAM node” on page 381](#)
- [“Rebuild of the sysplex-wide security associations structure \(EZBDVIPa\)” on page 381](#)
- [“Disconnect from the EZBDVIPa structure” on page 381](#)

Failure of a TCP/IP stack

If a TCP/IP stack (which is the distributor for a DRVIPA involved in IPSec connections) fails, ownership of the entries in the EZBDVIPa coupling facility structure associated with that DRVIPA (owned by the failing TCP/IP stack) is taken over. The ownership is taken over by the TCP/IP stack designated as the backup for the DRVIPA, if a backup stack exists. When the failing stack is restarted, ownership of those entries is given back to it.

Failure of a VTAM node

When a VTAM node fails, all TCP/IP stacks using that VTAM node to access the EZBDVIPa structure lose access to the structure, and the VTAM node is disconnected from the EZBDVIPa structure. When VTAM is restarted, the TCP/IP stacks that were using that VTAM node re-register as EZBDVIPa users with this VTAM and VTAM attempts to reconnect to the EZBDVIPa structure. When the reconnect is complete, VTAM notifies the TCP/IP users, and the TCP/IP stacks repopulate the EZBDVIPa structure with updates to the IPSec security associations that have changed while VTAM was down.

Rebuild of the sysplex-wide security associations structure (EZBDVIPa)

Rebuild can occur for several reasons:

- Operator initiated
- Loss of connectivity
- Structure failure

After a rebuild, the EZBDVIPa structure is repopulated using local data from each of the TCP/IP stacks. Each TCP/IP stack is notified that a rebuild has occurred, and the structure must be repopulated. Each TCP/IP stack sends updated information about the security associations for DRVIPAs that it owns to VTAM to be written to the EZBDVIPa structure.

During a rebuild, TCP/IP is not able to access the coupling facility. This creates a problem with obtaining sequence numbers. If TCP/IP cannot access the coupling facility, it cannot get the next available sequence numbers. Therefore, during rebuild, packets might be dropped and normal TCP/IP retransmit protocol will resend the data. It is assumed that rebuild takes on the order of 10-15 seconds, and the data is sent with the correct sequence number when the rebuild and repopulate have completed, so the connection is not broken. However, there are many factors which can affect the time for rebuild to complete, such as system load and number of structures being rebuilt. It is possible that connections might be dropped during such a rebuild. The situation is not likely, and it is rare that such a massive rebuild is required.

Disconnect from the EZBDVIPa structure

The following are cases in which VTAM might disconnect from the structure:

- A VARY NET,CFS,ACTION=DISCONNECT operator command.
- A rebuild fails and the original structure is not accessible.

- An internal error within VTAM occurs.

Disconnect is viewed as a much longer loss of coupling facility access than rebuild. It is more likely that connections will be dropped during the first case. However, VTAM has automatic reconnect logic for cases 2 and 3 above, which might occur before the connections timeout.

When VTAM reconnects, it notifies each TCP/IP stack using the EZBDVIP structure that a reconnect has occurred, and the structure must be repopulated. Each TCP/IP stack sends updated information about the security associations for DRVIPAs that it owns to VTAM to be written to the EZBDVIP structure.

Modifying the number of lists

To increase the number of lists in the EZBDVIP structure if you are not using subplexing within your sysplex, take the following steps:

1. Use the CFSIZER tool to determine the size and list count recommendations for the TCP/IP SWSA structure, EZBDVIP. See [“Determining the size of the coupling facility structure” on page 346](#).
2. Modify and install your CFRM policy with any recommended changes for INITSIZE and SIZE.
3. Issue the **MODIFY VTAMOPTS** command to modify the DVLSTCNT START option for all VTAMs in the sysplex. See [z/OS Communications Server: SNA Operation](#) for more information about the Modify command.
4. Issue the **SETXCF START,REBUILD** command for the EZBDVIP structure to rebuild the structure. See [SETXCF command](#) in *z/OS MVS System Commands* for more information.
5. Issue the **D NET,STATS,TYPE=CFS** command to verify the number of lists in the EZBDVIP structure.

If you are using subplexing within your sysplex, an increase in the number of lists for an EZBDVIP structure can affect other EZBDVIP structures that are defined in the VTAM subplex. The number of lists should be the same for all EZBDVIP structures in a VTAM subplex. Take the following steps to increase the number of lists in the EZBDVIP structures:

1. Use the CFSIZER tool to determine the size and list count recommendations for the TCP/IP SWSA structure, EZBDVIP, for each TCP/IP subplex in the VTAM subplex. If you have multiple EZBDVIP structures across the VTAM subplex, identify the largest recommended values. Use the largest recommended size and list count values for all EZBDVIP structures within the VTAM subplex. See [“Determining the size of the coupling facility structure” on page 346](#).
2. Modify and install your CFRM policy for all EZBDVIP structures with any recommended changes for INITSIZE and SIZE. The INITSIZE and SIZE values should be the same for all EZBDVIP structures within the VTAM subplex.
3. Issue the **MODIFY VTAMOPTS** command to modify the DVLSTCNT START option for all VTAMs in the VTAM subplex. See [z/OS Communications Server: SNA Operation](#) for more information about the Modify command.
4. Issue the **SETXCF START,REBUILD** command for each EZBDVIP structure in the VTAM subplex to rebuild the structure. See [SETXCF command](#) in *z/OS MVS System Commands* for more information.
5. Issue the **D NET,STATS,TYPE=CFS** command to verify the number of lists in each EZBDVIP structure. See [z/OS Communications Server: SNA Operation](#) for more information about the Display command.

Restriction: All VTAMs in a subplex must be at z/OS V2R3 or later to increase the number of lists in the EZBDVIP structure.

Tip: If the rebuild fails, the EZBDVIP structure might not be large enough. After the rebuild failure, there might not be a connection from VTAM to the EZBDVIP structure anymore. Review the CFSIZER recommendation for the size of the EZBDVIP structure and ensure that a policy that implements the recommendation is installed. After verifying that the structure size is correctly configured and installing the policy, you can use the following command at each VTAM to reestablish VTAM's connection to the EZBDVIP structure: **V NET,CFS,ACTION=CONNECT,STRNAME=EZBDVIP**.

Verify that the correct number of lists is in use by the structure by using the **D NET,STATS,TYPE=CFS** command.

Sysplexports

TCP/IP uses the coupling facility in the MVS sysplex to provide coordinated sharing of ephemeral port assignments among TCP/IP stacks in the sysplex, for Distributed DVIPAs. The configuration and use of this function is described in the [z/OS Communications Server: IP Configuration Guide](#).

Coupling facility failures for Sysplexports

There are four types of events that can impact the Sysplexports function:

- Failure of a TCP/IP stack
- Failure of a VTAM node
- Rebuild of the Sysplexports structure EZBEPORT
- Disconnect from the EZBEPORT structure

Failure of a TCP/IP stack

If a TCP/IP stack fails, and that stack is a target stack for a DRVIPA using the Sysplexports function, the distributor stack for that DRVIPA returns all ephemeral ports allocated to the failing stack back to the available pool in the EZBEPORT structure. If the distributor stack for a DRVIPA fails, and has no backup stack designated, the target stacks for that DRVIPA return all ephemeral ports assigned to each target stack back to the available pool. The ephemeral ports assigned to a failing distributor stack (with no backup) are recovered when the distributor stack is restarted and requests an ephemeral port for the first time.

Failure of a VTAM node

When a VTAM node fails, all TCP/IP stacks using that VTAM node to access the EZBEPORT structure lose access to the structure, and the VTAM node is disconnected from the EZBEPORT structure. When VTAM is restarted, the TCP/IP stacks that were using that VTAM node re-register as EZBEPORT users with this VTAM, and VTAM attempts to reconnect to the EZBEPORT structure. When the reconnect is complete, VTAM notifies the TCP/IP users, and the TCP/IP stacks repopulate the EZBEPORT structure with the list of ephemeral ports each TCP/IP stack has currently allocated.

Rebuild of the Sysplexports structure (EZBEPORT)

Rebuild can occur for several reasons:

- Operator initiated
- Loss of connectivity
- Structure failure

After a rebuild, the EZBEPORT structure is repopulated using local data from each of the TCP/IP stacks. Each TCP/IP stack is notified that a rebuild has occurred, and the structure must be repopulated. Each TCP/IP stack sends updated information about the ephemeral ports VTAM allocated to it to VTAM to be written to the EZBEPORT structure. During a rebuild, TCP/IP is not able to access the coupling facility. This creates a problem with obtaining ephemeral port numbers for new connections. Therefore, requests for connections involving new ephemeral port numbers fail during a rebuild. It is assumed that rebuild takes on the order of 10-15 seconds, and the requests for new ephemeral port numbers are satisfied when the rebuild and repopulate has completed. However, there are many factors which can affect the time for rebuild to complete, such as system load and number of structures being rebuilt. This situation is not likely, and it is rare that such a massive rebuild is required.

Disconnect from the EZBEPOR structure

The following are cases in which VTAM might disconnect from the structure :

- A VARY NET,CFS,ACTION=DISCONNECT operator command.
- A rebuild fails and the original structure is not accessible.
- An internal error within VTAM occurs.

Disconnect is viewed as a much longer loss of coupling facility access than rebuild. When VTAM reconnects, it notifies each TCP/IP stack using the EZBEPOR structure that a reconnect has occurred, and the structure must be repopulated. Each TCP/IP stack sends updated information about the ephemeral ports VTAM has allocated to it to VTAM to be written to the EZBEPOR structure.

Chapter 15. Implementing an APPN network

An APPN network contains network nodes and end nodes connected by APPN connections. VTAM can be coded as either a network node or an end node. Configuring VTAM for APPN is optional. To configure VTAM as an APPN node, select the characteristics of the node and define the resources of that node to VTAM.

The HOSTSA, NODETYPE, and SACONNS start options are used to determine what type of node VTAM will be. Table 46 on page 385 shows how these start options are defined for the different types of nodes.

Table 46. Start options and node type relationship			
Node type	NODETYPE	HOSTSA	SACONNS
Subarea node	not coded	<i>subarea number</i>	YES
Interchange node	NN	<i>subarea number</i>	YES
Migration data host	EN	<i>subarea number</i>	YES
Network node	NN	<i>subarea number</i> *	NO
End node	EN	<i>subarea number</i> *	NO
*HOSTSA is optional. If not specified, VTAM will start as a EN or NN. If HOSTSA is specified, then SACONNS=NO must be specified to start VTAM as an EN or NN.			

By coding the NODETYPE start option, you specify a node as an APPN node. After NODETYPE is coded, the following start options default to the values needed to enable APPN connections.

CPCP

The CPCP start option defaults to LEASED and enables CP-CP control sessions across APPN leased links. The other possible values for CPCP are SWITCHED, YES (for all links), and NO (for no links). For information about modifying CP-CP session enablement, see [“Enabling control sessions” on page 387](#).

CONNTYPE

The CONNTYPE start option defaults to the value APPN and enables full APPN function across a link. The other connection type value is LEN. There is also a CONNTYPE operand on the PU definition statement that can be used to override the value of the start option for specific links. For more information about the CONNTYPE start option and operand, see [Chapter 16, “Implementing a combined APPN and subarea network,” on page 401](#).

Coding considerations for APPN resources

LU definitions are not needed if you use dynamic independent LUs. See [“Defining independent LUs” on page 190](#).

Coding examples for each of the connection types are included in this section following the general requirements. See [“Channel connections between APPN nodes” on page 40](#) and [“Leased connections between APPN nodes” on page 47](#).

If you take advantage of the CPCP and CONNTYPE start option defaults, you only need to consider the following requirements for an APPN connection:

- To enable CP-CP sessions over a switched APPN connection, the CPCP start option should be defined as YES or SWITCHED.
- The PU definition for the NCP in a composite network node or for the VTAM in a network node or end node needs to indicate exchange identification support (XID=YES). (For a switched connection, XID cannot be coded but is automatically set to YES.)

- Use different names for the CP name and the PU name of a node.
- If you choose to allow dynamic definition of your independent logical units by defining the DYNLU start option or operand as YES, you only need to define the independent LUs on their owning CP.
- If you choose not to allow dynamic definition of your independent logical units, each network node that needs to connect to the independent logical unit needs a CDRSC definition (or a model CDRSC definition) for the resource. If you chose to code the ALSLIST operand on the CDRSC statement, the ALSLIST operand should specify either the name of an adjacent link station that can be used to connect to the LU or to ISTAPNPU. ISTAPNPU is a generic representation for an APPN adjacent link station. If ISTAPNPU is included in the ALSLIST, VTAM uses any available APPN link station when calculating a route to the LU.

If you chose not to code ALSLIST, VTAM searches for the LU in the APPN network and, if found, adds ISTAPNPU to the ALSLIST dynamically.

Also, unlike a LEN adjacent link station, an APPN adjacent link station is not required to have a NETID that matches the NETID of the independent LU at session establishment.

- To establish a statically defined switched APPN connection, either the CPNAME operand or the IDBLK and IDNUM operands must be defined for each switched PU definition. If the CPNAME operand is defined, its value must be the same as the CP name of the node (the SSCPNAME start option for VTAM nodes) that owns the PU. If the values are not the same, VTAM cannot locate the PU and LU definitions. There may be more than one PU related to that same CP.

To establish a statically defined switched LEN connection, either the CPNAME operand or the IDBLK and IDNUM operands must be defined for each switched PU definition. If the CPNAME operand is defined, its value specifically identifies a specific PU and cannot be duplicated.

- To establish a dynamically defined switched connection, you can allow the DYNADJCP start option to default to YES and code the DYNPU operand on the GROUP definition statement in the XCA or NCP major node.
- You need to know the maximum number of sessions that will be established through each NCP in a composite network node. For NCP type 2.1 connections, this number is defined during NCP generation and cannot dynamically expand. However, for a direct VTAM connection, the session control blocks are held in the BSBUF buffer pool that can expand and contract as necessary.
- For an NCP connection, you should also know the maximum number of sessions that any particular type 2.1 node can establish per link station. If you do not specify this limit, a single independent LU in a type 2.1 node can exhaust all of the session control blocks generated in the NCP.

There is no corresponding control for type 2.1 nodes directly attached to a VTAM. However, because this is a buffer pool, it can expand as needed to accommodate more sessions.

- When an independent LU in a type 2.1 node acts as the primary logical unit, NCP must assign another network address to the PLU. This address pool is generated in NCP. Because this is a predefined pool, you should know the maximum number of PLU sessions that all type 2.1 node independent LUs attached to that NCP can establish concurrently. For a direct VTAM connection, the addresses are assigned from the element pool that is managed by VTAM.
- Depending upon the specific type of APPN connection and the type of node at each endpoint, High-Performance Routing (HPR) can be used. If required, HPR support can be disabled on a node-wide basis using the HPR=NONE start option, or limited using the HPR=ANR start option. Within a node providing HPR support, HPR can be disabled on individual links by using the HPR operand on the GROUP, LINE, PU or CDRM definition statements. See [“Limiting HPR support” on page 391](#) for additional information about the effects of the HPR start option. For further information about HPR, see [“High-Performance Routing \(HPR\)” on page 388](#).

Maximum APPN Locate size considerations

Although VTAM supports sending and receiving APPN Locate searches with a size up to a maximum of 16 KB, other products that act as APPN network nodes or end nodes might support much smaller maximum APPN Locate sizes (as small as 1 KB). This can cause problems for ENs that have many links (described by

TG Vectors) to other nodes in the network, because these TG Vectors often must be included on the APPN Locate search in order to compute the best session path.

If any such APPN products are present along the path of an APPN Locate search, and the APPN Locate search that is to be sent exceeds the maximum Locate size allowed by these products, then some of these TG Vectors will be removed from the APPN Locate search in order to prevent the maximum Locate size from being exceeded. (VTAM typically copies the TG Vectors for active links to the APPN Locate search in the order these TGs were activated, which means the most recently activated TGs are removed first.) When this occurs, the network node server responsible for computing the session path might not choose the best possible path or the route computation could fail entirely (if the TG Vector describing the best or only path was removed). To avoid this problem, whenever possible, minimize the number of links (TGs) that end nodes have to other nodes. Using a connection network to reduce the number of EN TG Vectors can also be helpful.

Enabling control sessions

The CPCP start option specifies whether an APPN node supports CP-CP sessions with an adjacent node. The CPCP start option is valid only for nodes at which the NODETYPE start option is specified. APPN nodes use CP-CP sessions to communicate topology, routing, directory, and session information to adjacent control points. If you use the default, CP-CP sessions are supported on all APPN leased links. Because CP-CP sessions carry network control information, they should be active as long as the node is active. Using the default allows you to control whether your APPN switched links support CP-CP sessions on a link-by-link basis by defining the CPCP operand.

You can use the CPCP start option to specify support for CP-CP sessions on leased links only, on switched links only, on both leased and switched, or on neither leased nor switched.

The CPCP operand can be coded on the PU definition statement for the local SNA, model, NCP, and switched major nodes.

CP-CP sessions between two VTAM nodes

CP-CP sessions between two VTAM nodes require both CDRSC and ADJCP definitions at each VTAM to represent the partner CP. The CDRSC definition represents the partner LU, and the ADJCP definition represents the adjacent CP at the connection level. ADJCP definitions can be created dynamically or predefined. When needed, CDRSC definitions are dynamically created for adjacent CPs, regardless of the value coded for DYNLU.

Defining adjacent APPN nodes

The DYNADJCP start option controls whether adjacent CP minor nodes can be created dynamically. An adjacent CP major node, ISTADJCP, is automatically created when a VTAM node is activated. An adjacent CP minor node describes the CP name, the network identifier, node type (end node or network node), and whether dynamic definition is allowed for LUs owned by this adjacent node.

When the NODETYPE start option is specified, the default for the DYNADJCP start option is DYNADJCP=YES. If you allow the DYNADJCP start option to default to YES, adjacent CP minor nodes are created as needed to provide control and management of connections to adjacent APPN nodes.

Restriction:

- If an HPR pipe is to be used, the DYNADJCP option is ignored.
- When you define a PU with CPNAME, the ADJACENT CP will automatically be built and stored in ISTADJCP when the PU is activated, even when it's still in CONCT status. When the first link to the CPNAME activates, DYNADJCP=NO will not take effect for pipes towards that CPNAME. With DYNADJCP=NO, links from 'unknown' CPNAMEs require a VBUILD TYPE=ADJCP major node which defines every CPNAME you want links with.

Note: Unless CDRSCs or appropriate model CDRSCs are predefined for adjacent CPs, CDRDYN=YES is also required for the dynamic creation of adjacent CP minor nodes.

If you define the DYNADJCP start option as NO, you need to define every potential adjacent CP within adjacent CP major and minor nodes. Connections are established with only those nodes you specify. If you do not allow dynamic ADJCPs, and you do not define any adjacent CPs with the ADJCP definition statement, all connection attempts will fail.

You can choose to allow dynamic ADJCP minor nodes to be created for some links and not allow them for others. Be aware that for each link where you restrict dynamic definition, you need to define each CP to which you will allow connections with an ADJCP definition statement.

If you want dynamic definitions on most links, you can define the DYNADJCP start option as YES, and override it by defining the DYNADJCP operand as NO for the links on which you want to restrict dynamic definitions.

If you want limited dynamic definitions, you can code the DYNADJCP start option as NO and override it by coding the DYNADJCP operand as YES for links on which you want dynamic definitions.

Defining the logon mode for CP-CP sessions

CP-CP sessions always use the CPSVCMG logon mode. This logon mode is defined in the IBM-supplied default logon mode table, ISTINCLM. Logon mode CPSVCMG must be included in ISTINCLM in order to establish CP-CP sessions.

If no CPSVCMG logon mode is defined, CP-CP session activation will fail, and VTAM will issue the following message:

```
IST264I REQUIRED LOGMODE NAME CPSVCMG UNDEFINED.
```

In addition, message IST663I will display the error sense code X'08210002' (not valid mode name at CP).

High-Performance Routing (HPR)

Migrating NCP connections to APPN connections results in the creation of additional NCP control blocks, which can require additional storage space and cycles in the NCP. In addition, when a node or link along a route fails, sessions using that route fail.

Using the High-Performance Routing (HPR), you can migrate NCP connections to APPN connections without incurring the associated increase in storage and cycles. HPR uses a rapid transport protocol (RTP) connection to transport session traffic between session endpoints. HPR routes can also traverse an existing subarea network, as HPR support provides for the mapping of HPR routes over VR-based TGs.

In addition, if a node or link on an HPR route fails and an alternate HPR route exists between HPR session endpoints, HPR automatically switches routes and sessions are not disrupted. A method is also available that the operator can use to force an HPR path switch.

To use HPR over an NCP you must have, at minimum, NCP Version 7 Release 3 (and 37xx controllers supported by NCP Version 7 Release 3). To use HPR over a subnetwork boundary where one or both of the border nodes are composite network nodes, you must have, at minimum, NCP Version 7 Release 5.

What is High-Performance Routing?

High-Performance Routing is a set of enhancements for APPN that meets the following requirements:

- Improved APPN data routing: HPR transports data at very high speeds by using low-level intermediate routing and by minimizing the number of flows over the links for error recovery and flow control protocols. The flows are minimized by performing these functions at the endpoints rather than at each hop (link) along the path.
- Improved APPN reliability: HPR switches paths within the HPR portion of the network to bypass link and node failures if an acceptable alternate path is available. This occurs transparently to the sessions; in other words, the session is not disrupted.
- Functional equivalence: HPR maintains functional equivalence with APPN. To do so, HPR continues to support priority routing, connection networks and multiple network connectivity. Priority routing allows

the capability for higher priority traffic to pass lower-priority traffic in intermediate nodes within the HPR portions of the network. HPR also routes across connection networks or subnetwork boundaries in much the same way as APPN.

HPR routes are not given preferential treatment by the APPN routing algorithm. Existing non-HPR APPN routes will also be used if they meet the requirements of the APPN Class of Service.

- Seamless migration: HPR is designed for *drop-in* migration. A given APPN node can be upgraded to the HPR level of function without taking down the network, without configuring new parameters at its adjacent nodes, and without any logistical complications or coordination.

What is Rapid Transport Protocol?

Rapid Transport Protocol (RTP) is a transport layer protocol that provides the following benefits:

- Flow control.
- End-to-end error recovery.
- Selective retransmission of lost packets.
- Nondisruptive rerouting of sessions.
- Segmenting and reassembly of packets.
- Translation of APPN to HPR protocols (boundary function) for sessions that have both HPR and non-HPR portions.
- Translation of subarea to HPR protocols for sessions that have subarea and HPR portions. This translation is limited to routes that have a:
 - Non-HPR VR-based TG adjacent to an HPR route.
 - Subarea portion adjacent to a one-hop HPR route where the HPR endpoints are an interchange node and an adjacent APPN node. (Some restrictions apply; for details, see [“Sessions traversing subarea and APPN networks”](#) on page 397.)

RTP uses an adaptive rate-based (ARB) congestion control algorithm that makes efficient use of network resources by providing a congestion avoidance and control mechanism. The basic approach used in this algorithm is to regulate the input traffic during changing network conditions. When the algorithm detects that the network is approaching congestion, it reduces the input traffic rate until these indications go away. The algorithm also allows more traffic to enter the network without exceeding the rate the receiver can handle.

At LU-LU session establishment, if a node providing the RTP function determines that at least some portion of the session route traverses one or more HPR-capable links and eventually terminates in another RTP node, it will establish an *RTP connection* between the two RTP nodes. Subsequent sessions can reuse existing RTP connections if the HPR portion of the route, the APPN Class of Service, and the network connection endpoints (NCEs) are the same. CP-CP sessions over APPN node-to-node connections also use RTP connections.

Tip: You can use the HPRSESLM start option or the MODIFY VTAMOPTS,HPRSESLM= command (if HPRSESLM is not currently set to DISABLED) to limit the number of LU-LU sessions assigned to each RTP connection. Limiting the number of sessions on a single RTP connection can result in performance improvement by allowing concurrent traffic on multiple RTP connections.

When a VTAM establishes a logical connection between itself and another RTP node, VTAM creates a dynamic PU to represent the connection. These PUs are easily identified because their names are always in the form CNRxxxxx (unless the DYNHPPFX start option has been used to change this default prefix).

Note: A model PU definition can be created to customize the characteristics of dynamically created PUs. Use the DYNTYPE=RTP operand for the model PU definition in the model major node.

RTP transmits data in a form known as a network layer packet (NLP). The NLP begins with a header that specifies the automatic network routing (ANR) information necessary for routing the packet through the network.

Note: The maximum data transmission size referenced by VTAM to set the maximum size of data transmitted is calculated from the maximum packet size returned by the Route_Setup signal exchanged during RTP initialization.

To prevent excessive data segmenting within VTAM's data transmission processing, carefully evaluate the TCP/IP MTU size (for Enterprise Extender) and the maximum received data size for VTAM and each node in the network.

For Enterprise Extender, VTAMs maximum BTU length the sender can receive is calculated as follows:

```
Max_Receive_Size=  
MAX((MIN(TCP/IP MTU SIZE -3),32767),768)
```

Because excessive segmenting can occur in the effort to conform to maximum data size requirements, you should optimize the maximum received data size. Failure to do so can result in the following situations:

- Inefficient TIPAC (I/O) buffer allocation
- Excessive TIPAC buffer pool expand operations
- Degraded data transmission performance

What is automatic network routing?

Automatic network routing is a low-level routing mechanism that minimizes cycles and storage requirements for routing packets through intermediate nodes. Unlike APPN, intermediate ANR nodes are not aware of SNA sessions or RTP connections passing through the node. All an ANR node must do is read the header in a network layer packet and forward the information to the next node on the path. The ANR information is learned by RTP during establishment of the RTP connection by sending a "Route Setup" message which flows through all nodes on the prospective HPR path.

Tip: On nodes configured as border nodes (BN=YES is specified) with HPR fully enabled (HPR=RTP is specified), code RTPONLY=YES on ADJCP definition statements in adjacent CP major nodes if you want the border node to maintain awareness of all sessions established to, from or through the adjacent node being defined. The RTPONLY operand is valid only when the activating node is configured as a border node, and RTPONLY=YES is meaningful only when the adjacent node being defined is a nonnative border node or nonnative network node. By coding RTPONLY=YES, you are instructing VTAM to disallow use of the ANR function for any new RTPs that are established or path switched to, from, or through the adjacent nonnative node being defined. (Allowing the use of ANR could result in RTPs being established through this border node, thereby preventing the border node from maintaining awareness of any sessions that use these RTPs.) Instead of allowing RTPs to be established through this border node, coding RTPONLY=YES will result in VTAM forcing these RTPs to terminate on this border node or forcing the use of intermediate session routing (ISR) instead of HPR (or a combination of the two).

Restrictions:

- Use of RTPONLY=YES at any APPN subnetwork boundary on a session setup path prevents the use of global VRNs (GVRNs) for intersubnetwork connectivity, because using GVRNs could result in sessions being established across this subnetwork boundary without this border node maintaining awareness of these sessions.
- Use of RTPONLY=YES can result in an increase in network traffic in the form of additional Route_Setup flows used for RTP establishment. These additional Route_Setup flows will occur only during the establishment of sessions that cross subnetwork boundaries defined with RTPONLY=YES.
- Use of RTPONLY=YES can result in an increase in storage and CPU utilization, because of VTAM maintaining awareness of these sessions and performing ISR instead of HPR/ANR for these sessions.

How does HPR switch paths?

When a link or node outage occurs in the path of the RTP connection, the RTP endpoint attempts to find an alternate route to maintain the connection. If an acceptable route is not immediately found,

RTP will continue to try to find one until the path switch timer expires. (The path switch timer duration is modifiable. See the HPRPST start option in [z/OS Communications Server: SNA Resource Definition Reference](#) for more details on the path switch timer.) While the path switch attempt is in progress, data transmission is suspended (because of the outage in the original path) until the path switch completes.

Additionally, a path switch attempt can be forced by the operator by invoking the MODIFY RTP command. This will force the specified RTP to make one attempt at finding a better route (as specified by the requirements of the Class of Service) for the RTP connection. If no better route is available, the original route for the RTP connection will be unchanged.

The PSRETRY start option can be used to force VTAM to perform the following tasks:

- Periodically attempt to find a better route for existing HPR connections
- Immediately attempt to find a better route when a link on this host is activated or changes status

Like MODIFY RTP processing, if a better route is not available, the original route for the connection is unchanged. For both timer-driven and operator-initiated path switch processing, a path switch can be performed to what appears to be an identical path if the path uses a VR-based transmission group. The reason is that VTAM will perform virtual route selection whenever the first hop in the path is a VR-based transmission group, and this underlying virtual route might change from one usage of the VR-based transmission group to the next. VTAM will not verify that a different underlying virtual route is selected, but instead will just perform the path switch in case a change occurred. For more information, see [z/OS Communications Server: SNA Resource Definition Reference](#).

The HPRPSMSG start option can be useful for large installations that have hundreds or thousands of RTP endpoints on z/OS Communications Server. This start option controls the HPR path switch message reduction function. In the event of a large network outage, this function can prevent VTAM from flooding the system console log with HPR path switch messages. Although this function reduces the amount of path switch information issued to the console log, the summary information attempts to supply the operator with enough information to understand the size and scope of the outage. For more information about the HPRPSMSG start option, see [z/OS Communications Server: SNA Resource Definition Reference](#).

HPR implementation overview

The default level of HPR support available is dependent upon:

- Node type (NODETYPE and HOSTSA start options)
- Connection type

Levels of HPR support

There are three levels of HPR support for a VTAM node:

Rapid Transport Protocol (RTP)

An RTP connection can exist between two VTAM nodes with APPN capability (NODETYPE is coded); these two VTAMs can be endpoints of an HPR route.

Automatic network routing (ANR)

A VTAM network node (NODETYPE=NN) can also provide ANR-level support as an intermediate node on an HPR route.

No HPR support

HPR support can be disabled for a particular VTAM network node or end node using the HPR=NONE start option. If HPR support is disabled, that VTAM cannot be an endpoint or intermediate node for an HPR route. A VTAM subarea node (NODETYPE not coded) also cannot provide HPR support.

Limiting HPR support

HPR support provided by a VTAM APPN node can be limited using the following start options:

HPR=ANR

You can use the HPR=ANR start option to limit HPR support provided by a network node. Ordinarily, network nodes can provide both RTP-level and ANR-level support, but use of HPR=ANR precludes a network node from providing RTP-level support. If you use HPR=ANR, coding HPR=YES on a CDRM or PU definition statement (or allowing it to default to YES) results in ANR-level HPR support over that connection.

HPR=(RTP,ANR)

Using HPR=(RTP,ANR) indicates this VTAM provides RTP-level HPR support only if HPR=YES is coded for a particular CDRM or PU definition statement. If you let the HPR operand default on a CDRM or PU definition statement, you only get ANR-level HPR support over that connection. This way you can easily select exactly which links serve as RTP endpoints, while allowing others to provide ANR-level HPR support by default.

Note: The HPR=(RTP,ANR) start option is recommended only for migration purposes. It can cause unexpected results in HPR route selection. The ANR part of this option means that the nodes connected to a TG can act only as intermediate nodes for an HPR connection if HPR=YES is not coded on the PU definition statement for that connection or the PU is not activated with HPR=YES on the VARY ACT command. They cannot act as endpoints for the connection.

If HPR=(RTP,ANR) is used, it is important to ensure that a TG is defined with the same capabilities on both sides of the link. To use the TG with ANR capabilities for the origin and destination nodes of that TG, ANR must be defaulted for the PUs defining the TG on both the origin and destination nodes, because there is no other way to specify ANR for a PU. To use a TG with RTP capabilities for the origin and destination nodes of the TG, with HPR=(RTP,ANR), HPR=YES must be specified on the PU statement that represents the link station, or the PU must be activated with HPR=YES on the VARY ACT command.

In the following example, TG15 and TG17 cannot be the first hop or the last hop of an HPR connection, but TG16 can. The combination of definitions for TG15, with HPR=YES on one side and ANR defaulted on the other, is not recommended.

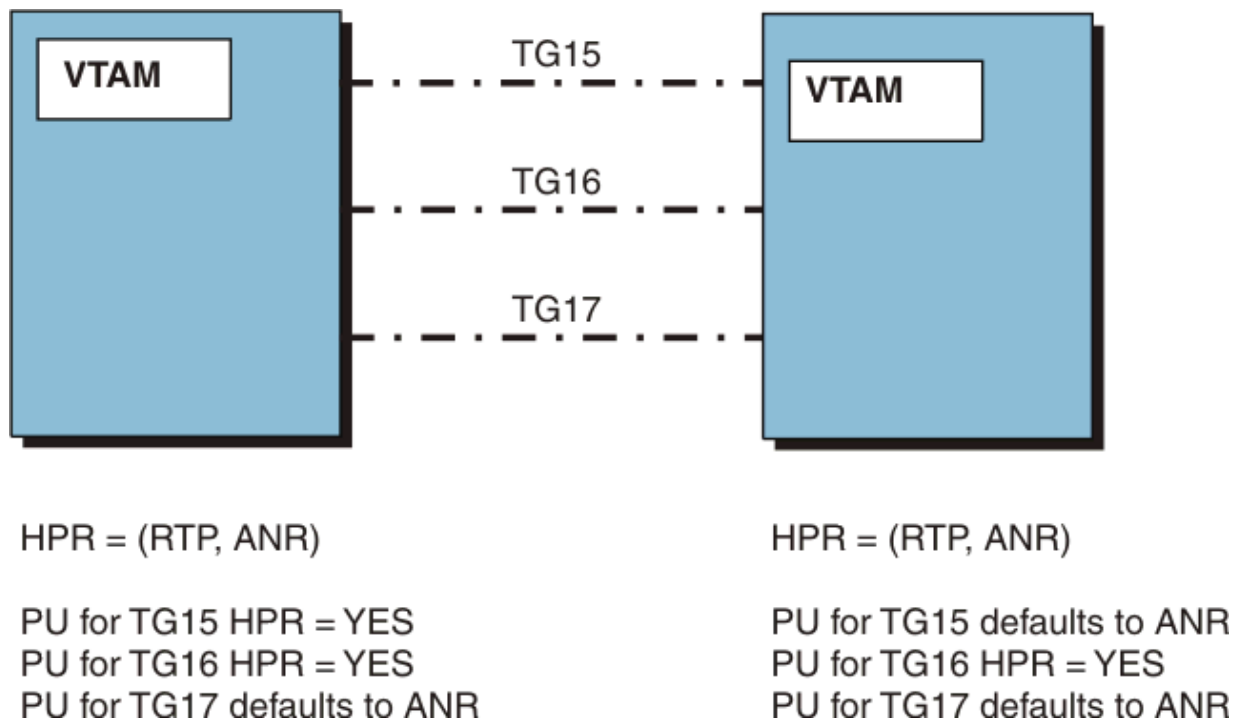


Figure 106. HPR=(RTP,ANR) and TG capabilities

Matching HPR characteristics is especially critical when defining TGs across a border node connection (ISL) and when defining connections to an end node (endpoint TGs). With the current APPN architecture, the HPR characteristics are known in only one direction, so the node doing HPR route

calculation cannot ensure that the nodes on both sides of the TG have the necessary HPR capabilities. In these cases, the HPR capabilities must be the same on both sides of the TG.

In summary, unless you adhere to the guidelines above, the establishment of HPR connections will not work correctly across border node TGs (ISLs) or TGs to an end node (endpoint TGs). In addition, TGs used for enterprise extender must be RTP capable on both sides.

HPR=(RTP,NONE)

Specifying HPR=(RTP,NONE) indicates this VTAM provides RTP-level HPR support only if HPR=YES is coded on the CDRM or PU definition statement. If you let the HPR operand default on a CDRM or PU definition statement, you get no HPR support over that connection.

HPR=(ANR,NONE)

Specifying HPR=(ANR,NONE) indicates this VTAM provides ANR-level HPR support only if HPR=YES is coded on the CDRM or PU definition statement. If you let the HPR operand default on a CDRM or PU definition statement, you get no HPR support over that connection. This way you can get ANR support on selected links without automatically getting it on all links.

HPR support can also be disabled by coding HPR=NO on a particular GROUP, LINE, or PU definition statement, or on the CDRM definition statement for VR-based TGs.

Sample HPR configurations

HPR over composite network nodes

In Figure 107 on page 393, HOSTB and HOSTC are composite network nodes connected by a VR-based TG (TG255). In this configuration, an LU-LU session between HOSTA and HOSTD can use an RTP connection that traverses the composite network nodes. HOSTB and HOSTC have no awareness of the session between HOSTA and HOSTD, as they are providing ANR-level support only. However, for sessions where HOSTB or HOSTC constitute one end of a session, HOSTB and HOSTC can provide RTP-level support. HOSTA, as a network node, can also provide ANR-level support for sessions between HOSTB, HOSTC, or HOSTD and some other unpictured node. HOSTD, as an end node, cannot provide ANR-level support.

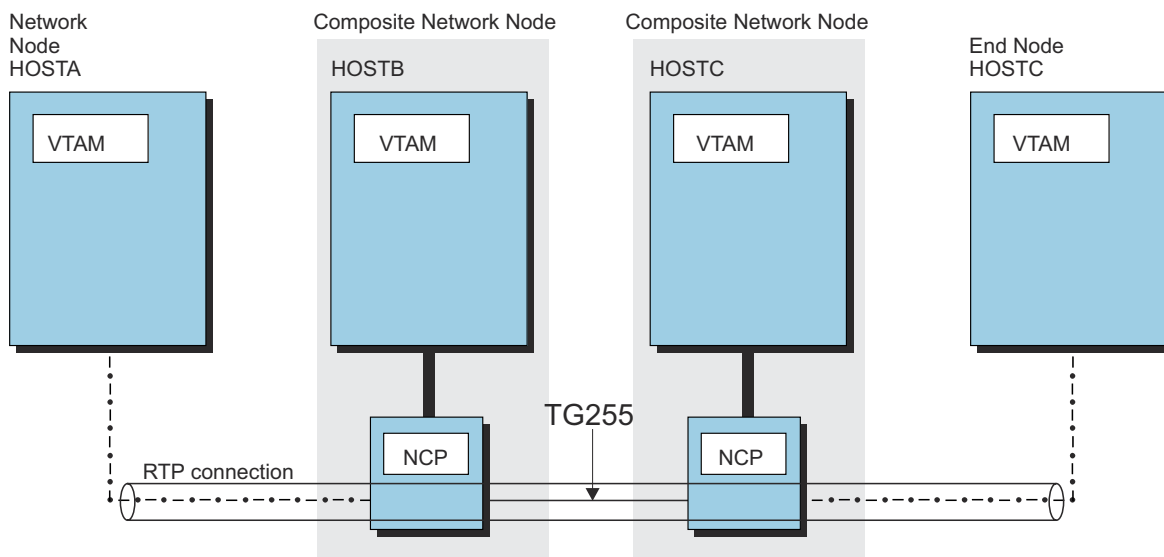


Figure 107. HPR over composite network nodes

HPR over APPN host-to-host channel connections

In Figure 108 on page 394, an LU-LU session between HOSTA and HOSTC can use an RTP connection that traverses the APPN host-to-host channel connections between HOSTA and HOSTB and between HOSTB and HOSTC. For sessions between HOSTA and HOSTC, HOSTB provides only ANR-level support. However, for sessions between HOSTB and HOSTC, HOSTB can provide RTP-level support as the endpoint of an

HPR route. Similarly, though not pictured, HOSTB can provide RTP-level support for sessions between HOSTB and HOSTA.

Note: Connections between two hosts that both provide RTP-level support can use HPDT MPC. See [“Multipath channel connections” on page 40](#) for information about the different levels of MPC provided by VTAM.

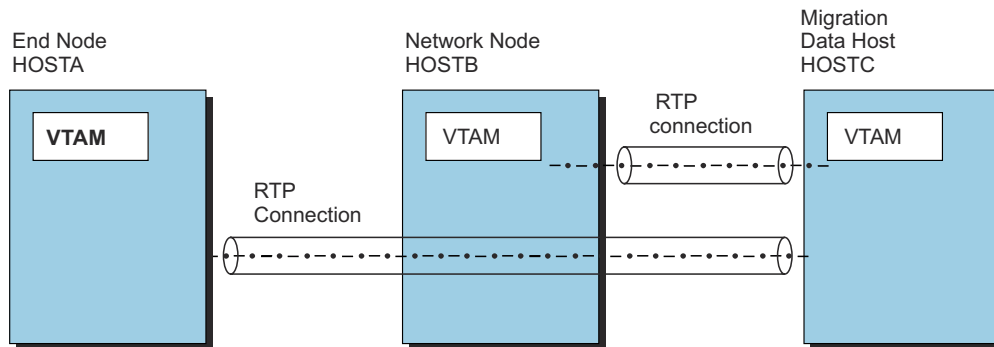


Figure 108. HPR Over APPN host-to-host channel connections

Multiple HPR routes between HPR session endpoints

In Figure 109 on page 395, like Figure 108 on page 394, none of the hosts have NCPs and all have APPN capability. In this configuration, an LU-LU session between HOSTA and HOSTD can use an RTP connection that traverses the APPN host-to-host channel connections between HOSTA and HOSTB and between HOSTB and HOSTD. For sessions between HOSTA and HOSTD, HOSTB provides only ANR-level support. However, for sessions between HOSTB and HOSTA or HOSTB and HOSTD, HOSTB can provide RTP-level support as the endpoint of an HPR route.




Figure 109. Multiple HPR routes between HPR session endpoints

Also in [Figure 109 on page 395](#), note that an alternate HPR route exists between HOSTA and HOSTD through HOSTC. When multiple paths exist between HPR session endpoints, HPR support provides for automatic, nondisruptive path switching when a node or link fails on the path in use. For example, if the APPN host-to-host channel connection between HOSTA and HOSTB fails, HPR support provides for the movement of the RTP connection to the alternate path between HOSTA and HOSTD through HOSTC. LU-LU sessions continue uninterrupted.

Multiple RTP connections over the same HPR route

Multiple RTP connections can exist over the same HPR route. For example, in [Figure 109 on page 395](#), multiple RTP connections can exist between HOSTA and HOSTD through HOSTB. Furthermore, it is possible that more than one of these connections are for sessions requiring the same Class of Service. In this case, to determine which connection a session is traversing, use the DISPLAY SESSIONS command at one of the RTP endpoints. First use DISPLAY SESSIONS to find the session identifier (SID) for the session in question. Then issue DISPLAY SESSIONS and specify the SID on the SID operand. The resulting display shows the RTP PU (CNRnnnnn, unless the DYNHPPFX start option has been used to change the default prefix) as the ALSNAME, if the session is over an HPR route.

Using VR-based TGs for non-HPR endpoints

As described previously, with HPR, you can migrate NCP connections to APPN connections without incurring the associated increase in storage and cycles. However, this is not always possible. In [Figure 110 on page 396](#), for instance, for a session between the terminal and HOSTB (an interchange node), HPR cannot be used. If APPN storage demands are unacceptable in such a situation, add a VR-based TG to connect the terminal and HOSTB through the subarea network, which allows them to act as APPN peers. Meanwhile, for sessions between HPR-capable endpoints, HPR routing can continue over the existing APPN connection.




Figure 110. Using VR-based TGs for non-HPR endpoints

Sessions involving HPR and APPN routes

Some sessions might have routes consisting of both APPN segments and RTP segments. In [Figure 111 on page 397](#), for example, the VTAM on HOSTC might not have been defined to provide HPR support. In this case, for a session between HOSTB and HOSTE, HPR is used where possible (between HOSTB and HOSTA and between HOSTD and HOSTE). Otherwise, APPN routing is used (between HOSTA and HOSTC and between HOSTC and HOSTD).




Figure 111. Session involving HPR and APPN routes

Sessions traversing subarea and APPN networks

When one end of a session is in a subarea network and the other in an APPN network, HPR routing can be used at the interchange node in the session path, but only under certain conditions.

APPN session endpoint is one hop from the interchange node

If the session is from a subarea node through an interchange node to an end node that is one hop away from the interchange node, translation of subarea to HPR can occur. In [Figure 112 on page 397](#), for example, a session can exist between LUA and LUB, and HOSTB will translate subarea data received from HOSTA to HPR and route it across the RTP connection to HOSTC.




Figure 112. Interchange node using HPR routing between subarea and APPN

Similarly, if the session is from a subarea node through an interchange node to a network node that is one hop away from the interchange node, translation of subarea to HPR can occur if the interchange node is IBM Communications Server for OS/390 V2R10 or later and the network node is one of the following types:

- A network node with IBM Communications Server for OS/390 V2R10 or later installed
- A non-z/OS Communications Server network node

For instance, in [Figure 112 on page 397](#), if HOSTC were a network node, HOSTB will translate subarea data received from HOSTA to HPR and route it across the RTP connection to HOSTC, provided that HOSTB and HOSTC meet the above criteria.

APPN session endpoint is more than one hop from the interchange node

If the session endpoint in the APPN network is an end node or network node that is more than one hop away from the interchange node, translation of subarea to HPR can occur if the interchange node is IBM Communications Server for OS/390 V2R10 or later and the adjacent network node is one of the following types:

- A network node with IBM Communications Server for OS/390 V2R10 or later installed
- A non-z/OS Communications Server network node

Note, however, that the RTP connection will only extend from the interchange node to the adjacent network node. An additional, separate, RTP connection can be established for the remainder of the APPN session path (depending on the HPR and RTP capabilities of the nodes on the path), but a single end-to-end RTP connection between the interchange node and the APPN session endpoint will not be used. In [Figure 113 on page 398](#), for example, HOSTD is part of a subarea network connected to an APPN network through HOSTC. In the case of a session between LUD and LUA, one RTP connection can be established between the HPR-capable endpoints HOSTC and HOSTB, and a second RTP connection can be established between HOSTB and HOSTA. Of course, for sessions between LUC and LUA, a single end-to-end RTP connection can be used between HOSTC and HOSTA (through HOSTB).

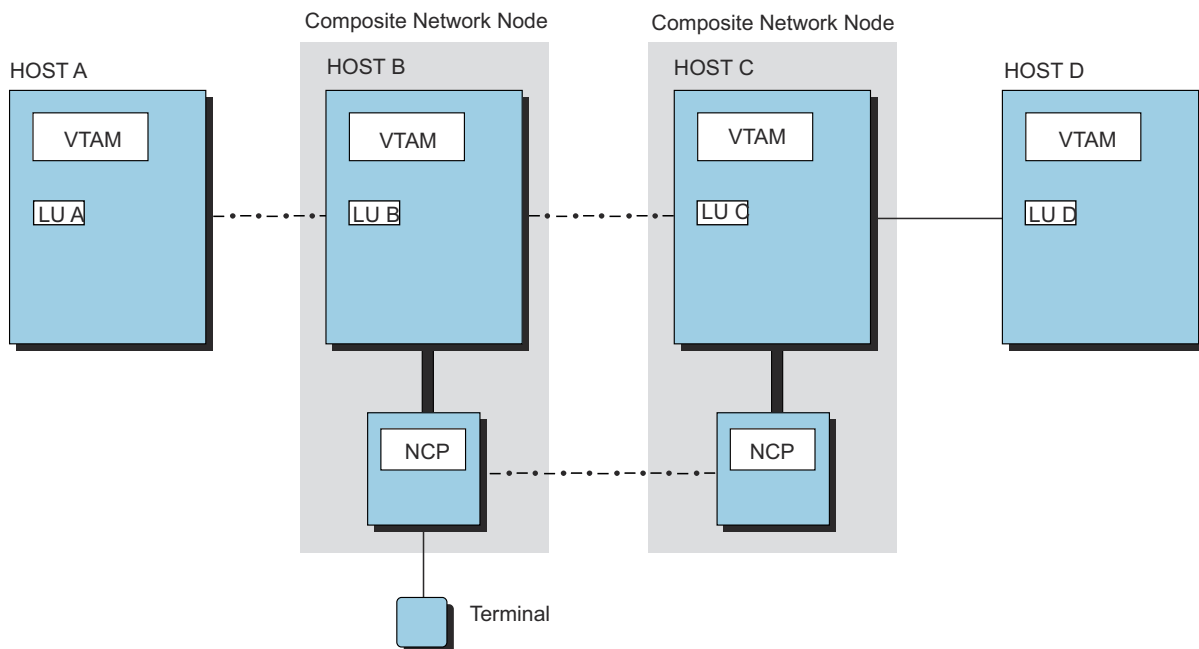


Figure 113. Sessions traversing APPN and subarea networks

Session endpoints are in separate subarea networks

When the session endpoints are in separate subarea networks connected by one or more APPN networks, translation of subarea to HPR can occur if the interchange nodes are IBM Communications Server for OS/390 V2R10 or later. For instance, in [Figure 113 on page 398](#), if HOSTA and HOSTD are both subarea nodes and there is a session between LUA and LUD through HOSTB and HOSTC, HOSTB and HOSTC will

translate subarea data received to HPR and route it across the RTP connection between them. If the interchange nodes are separated by other APPN nodes, the RTP connection from each interchange node is limited to its adjacent node. The RTP connection between intermediate nodes can include multiple nodes.

Setting session paths using HPRNCPBF

HPRNCPBF is a VTAM start option that you can use to determine the path used for session setup for LUs whose entry point to the HPR network is by way of NCP/FID4. This lets you determine whether HPR should be used in cases where it will cause session data to travel through an NCP twice.

In this case, a more optimal HPR TG between the VTAM nodes exists. When the LU requests a session with *appla*, the session is usually set up directly through the NCP to *appla* on VTAM2. While this path might offer better performance, it cannot take advantage of HPR features, such as dynamic rerouting or the high availability offered by MNPS. HPRNCPBF can be set to either YES or NO (the default). If you select NO, the session continues to be set up over the direct path. If you select YES, the session will be set up over the HPR pipe between the two VTAM nodes. Specifically, it will use the VR-TG path. In this case, in a connection failure between the NCP and VTAM, the session is redirected over an alternate route without disruption.

Note: If you select YES, the session requests from the NCP-attached LU are routed through the NCP twice.

For more information, see [z/OS Communications Server: SNA Resource Definition Reference](#).

Chapter 16. Implementing a combined APPN and subarea network

A sample communication management configuration (CMC) is shown in [Figure 114 on page 401](#). This configuration is not intended to be typical or recommended, but is used here to demonstrate converting to APPN.

Before using the APPN capabilities of VTAM, any communication controller that is to support APPN links (all of them in [Figure 114 on page 401](#)) must be a 3745 with NCP Version 6 Release 2 or later. Communication controllers that do not support APPN links do not require NCP Version 6 Release 2 or later.




Figure 114. Example of communication management configuration

Use the following steps to begin to convert this configuration to APPN:

1. Determine which connections to type 2.1 nodes you want to be LEN connections, and which connections you want to be APPN connections.

All connections between the NCPs and type 2.1 nodes are automatically tried as APPN connections when APPN support is specified, unless you explicitly designate them as LEN connections, using the `CONNTYPE=LEN` start option or the `CONNTYPE=LEN` operand on the `GROUP`, `LINE`, or `PU` definition statement. By specifying `CONNTYPE=LEN` initially, connections can later be converted to APPN in phases.

Notes:

- a. Prior versions of VTAM did not enforce CPNAME uniqueness for LEN nodes, except when CPNAME is coded on a switched PU definition statement. It is possible that your network has PUs with duplicate CPNAMEs. This CPNAME duplication should be resolved before a node becoming an APPN node. Specifying `CONNTYPE=LEN` for a connection causes VTAM to avoid checking for duplicate CPNAMEs, except when CPNAME is coded on a switched PU definition statement.

- b. When converting a connection from LEN to APPN, ensure that the adjacent link station name (VTAM-defined PU name) for an adjacent node is unique from the CP name of that node.

If you want most or all connections to type 2.1 nodes to initially be LEN connections, code the CONNTYPE=LEN start option at the CMC host. Then, if you do not code CONNTYPE=APPN on any GROUP, LINE, or PU definition statement, all connections to type 2.1 nodes are LEN connections.

If you want only particular connections to type 2.1 nodes to initially be LEN connections, do not use the CONNTYPE=LEN start option. Instead, code CONNTYPE=LEN on the appropriate GROUP, LINE, or PU definition statements.

2. Code the NODETYPE=NN start option at the CMC host.

The NODETYPE start option is required to specify APPN support. Assuming that the HOSTSA start option is also coded for HOSTC, coding NODETYPE=NN makes this VTAM an interchange network node. An interchange network node supports subarea and APPN protocols, and is capable of transforming subarea protocols to APPN protocols, and vice versa.

To APPN, the CMC host and all owned NCPs (NCPC1, NCPC2, and NCPC3) become one composite network node, as shown in [Figure 115 on page 402](#).

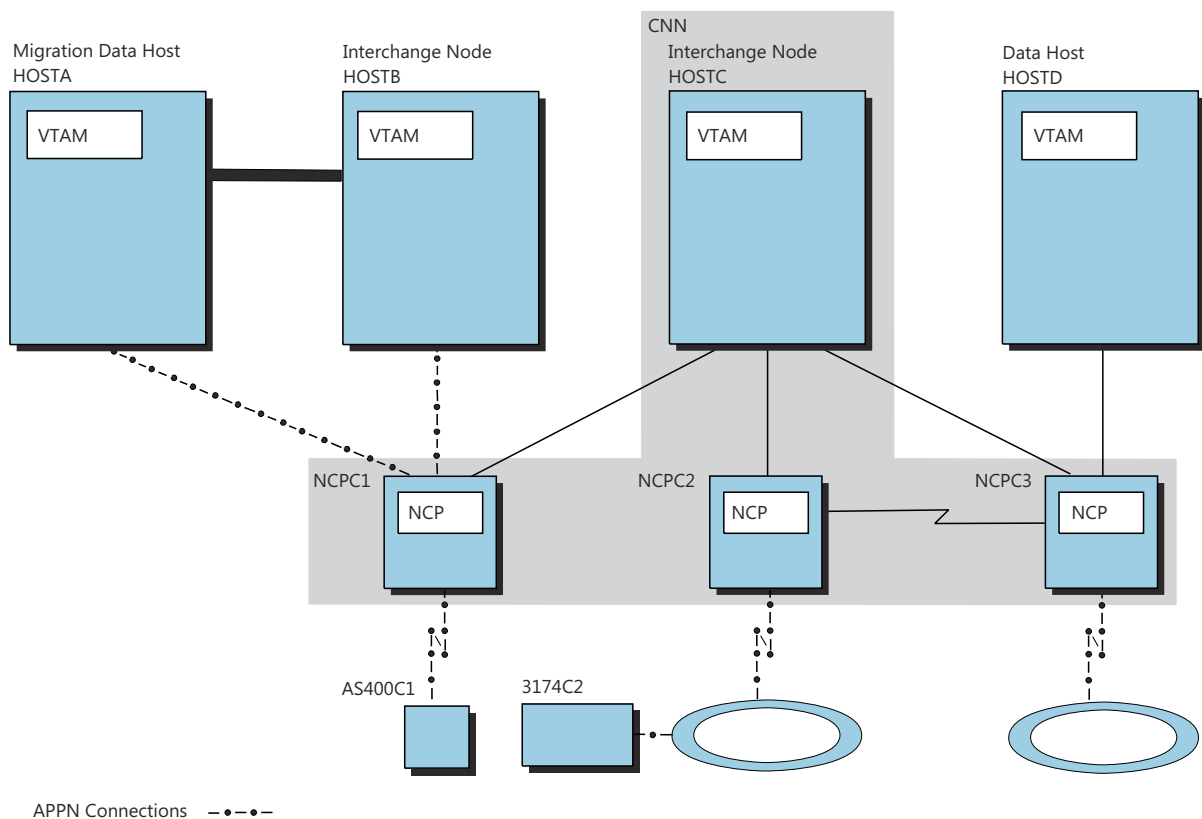


Figure 115. Communication management configuration after conversion

At this point, however, provided that all connections to type 2.1 nodes were specified as LEN connections in step “1” on page 401 using CONNTYPE=LEN, VTAM and NCP perform only their traditional subarea functions. If VTAM is started in HOSTC, no CP-CP sessions would be started between the composite network node and the adjacent nodes because connections to type 2.1 nodes are LEN connections and the adjacent VTAMs do not have APPN support.

3. Depending on your conversion strategy determined in step 1, convert LEN connections to APPN connections as required by doing one of the following actions:
 - Changing CONNTYPE=LEN to CONNTYPE=APPN (or deleting CONNTYPE=LEN) on the appropriate GROUP, LINE, or PU definition statements (and deactivating and reactivating the appropriate links), or restarting VTAM with the CONNTYPE=APPN start option.

- Specifying any other required APPN operands on the GROUP, LINE, or PU definition statement (for example, NN, CPCP, or TGP).

Note: If the NN operand is not specified, the node type (NN or EN) of the adjacent type 2.1 node is determined when the connection is activated. If specified incorrectly, the connection will fail.

- Defining the connection as APPN at the adjacent type 2.1 node.

In the example in [Figure 115 on page 402](#), all LEN connections are converted to APPN connections. With APPN connections, HOSTC can be used as the network node server, and independent LUs can now set up LU-LU sessions with the composite network node using APPN searches and flows. In addition, if a new independent LU is added to a type 2.1 node that is adjacent to the composite network node, the independent LU must now be defined only to that adjacent type 2.1 node. VTAM requires no definition for the added LU.

4. Code the NODETYPE=NN start option at data host B. Code a local SNA major node (VBUILD TYPE=LOCAL) that defines the type 2.1 channel attachment between data host B and NCPC1. Code PUTYPE=2 on the PU definition statement in the local SNA major node.

Assuming that the HOSTSA start option is also coded at HOSTB, this VTAM is now an interchange network node as shown in [Figure 115 on page 402](#).

To the composite network node, HOSTB appears as a network node. The connection between HOSTB and NCPC1 is now a type 2.1 connection. A CP-CP session is now possible between HOSTB and the composite network node. Thus, LUs connected to the composite network node through one of the NCPs can now set up LU-LU sessions with any application in HOSTB using APPN searches and flows. New applications added in HOSTB need not be defined to the composite network node.

The change of the link between HOSTB and NCPC1 to an APPN link results in the creation of additional NCP control blocks, which could require additional storage space. Whether you see a net decrease or increase in NCP storage is dependent upon the decrease in path table sizes in relation to any increase in storage because of the number of cross-domain sessions that are now boundary function sessions over APPN connections.

5. Code the NODETYPE=EN start option at data host A. Again, use a local SNA major node (VBUILD TYPE=LOCAL) to define the type 2.1 channel attachment between data host A and NCPC1, and code PUTYPE=2 on the PU definition statement in the local SNA major node.

Assuming that the HOSTSA start option is also coded at HOSTA, this VTAM is now a migration data host as shown in [Figure 115 on page 402](#). A migration data host node supports subarea and APPN protocols, but does not act as an intermediate node. A migration data host cannot activate any NCPs.

To the composite network node, HOSTA appears as an end node. The connection between HOSTA and NCPC1 is now a type 2.1 connection. A CP-CP session is now possible between HOSTA and the composite network node. Thus, LUs connected to the composite network node through one of the NCPs can now set up LU-LU sessions with any application in HOSTA using APPN searches and flows. New applications added in HOSTA need not be defined to the composite network node.

The change of the link between HOSTA and NCPC1 to an APPN link results in the creation of additional NCP control blocks, which could require additional storage space.

A migration data host can also still participate in SSCP-SSCP sessions with another VTAM (for example, HOSTB). If the channel-to-channel (CTC) connection between HOSTA and HOSTB is not also converted to APPN, you can continue to use subarea CTC support over that connection.

In this example, the NODETYPE start option is not coded at HOSTD. HOSTD can still communicate with resources in the HOSTC domain using the subarea connection to NCPC3. HOSTD can also still communicate with resources in the HOSTA and HOSTB domains using the interchange node function of HOSTC.

As a result of the above conversion steps, the following definitions are no longer required. You can remove them, but they do not have to be removed.

Note: If VR-based TGs are going to be used, you need to maintain definitions for routes and CDRMs. See [“Virtual-route-based transmission groups” on page 80](#) for additional information.

- All PATH definitions for routes between HOSTA or HOSTB and other subareas, except PATHs for the CTC connection between HOSTA and HOSTB
- In HOSTA and HOSTB, the CDRM definitions for HOSTC and HOSTD
- In HOSTC and HOSTD, the CDRM definitions for HOSTA and HOSTB
- CDRSC definitions in HOSTA and HOSTB for LUs in HOSTC and HOSTD, if present
- CDRSC definitions in HOSTC and HOSTD for LUs in HOSTA and HOSTB, if present
- Wildcard definitions in 3174C2 and AS400C1 that allow requests for unknown resources to be forwarded to VTAM over LEN links (assuming the links to 3174C2 and AS400C1 are now APPN links)
- ALS selection logic in the session management exit that allows requests to be forwarded from VTAM to independent LUs over LEN links that have been converted to APPN
- APPN-capable PUs included on the ALSLIST operand on the CDRSC definition statement (They can be optionally replaced by a single ISTAPNPU entry, a reserved keyword that represents any APPN capable PU.)

Start options defining a combined subarea and APPN environment

In addition to the CONNTYPE, NODETYPE, and HOSTSA start options, some other start options used in a combined subarea and APPN environment are associated with locating resources. All new and changed start options default to optimal values for an APPN network. During conversion, however, these defaults can be modified based on your particular conversion strategy and configuration.

SORDER start option

SORDER controls the order in which the APPN and subarea portions of the network are searched when a search request is processed by an interchange node or migration data host. SORDER can be specified as a start option or coded on individual ADJSSCP table definitions (or both). If SORDER is not coded for a given ADJSSCP table, the current value of the SORDER start option value is used each time that ADJSSCP table is selected for use.

Note: If SSCORD=PRIORITY is specified and previous searches for the target resource have been performed, the search order might be modified based on the results of these previous searches.

For further information about SORDER, see [“Controlling searches” on page 413](#).

SSEARCH start option

The setting of the SSEARCH start option determines whether the subarea network is searched when search requests from the APPN network arrive at an interchange node. By default, the subarea network is searched.

Note: Resources in the domain of the interchange node are found even when SSEARCH=NO is specified.

For more information about the SSEARCH start option, see [“Controlling searches” on page 413](#).

CDRSCTI start option

The setting of the CDRSCTI start option controls the amount of time that a dynamic CDRSC is retained by VTAM after its last session ends.

However, dynamic CDRSCs representing APPN resources are freed by an interchange node immediately after the last session ends, regardless of the specified CDRSCTI value. This is because APPN directory services in the interchange node has an entry for each APPN resource.

Dynamic CDRSCs representing subarea resources are retained for eight minutes (by default). You can increase the CDRSCTI start option value to minimize the potential for unnecessary broadcasts of the APPN network for subarea resources.

For further information about the CDRSCTI start option, see [“Controlling searches” on page 413](#).

IOPURGE start option

The setting of the IOPURGE start option specifies an interval after which outstanding I/O requests are purged. The types of I/O requests that are checked are:

- CDINIT requests
- Direct search list requests
- APPN search requests
- HPR route setup requests

For CDINIT requests, after the specified time interval, if a response has not been received, VTAM continues its search through the adjacent SSCP table until it either finds the resource or the table has been exhausted. For the other requests, if IOPURGE expires, VTAM handles the request as if a negative response to the request has occurred.

If IOPURGE is set too low, outstanding searches will be terminated before normal searching has completed. This can result in failed searches and possibly increased network traffic because of retries. This can occur if IOPURGE on an end node or migration data host is set to a value that does not allow the network node server to complete all search tasks and return a reply.

If IOPURGE is set to 0, outstanding I/O requests can remain outstanding indefinitely.

Note: Do not use IOPURGE for congestion control. If you have APPN congestion control problems, use a lower MAXLOCAT value. IOPURGE was intended for situations where a node drops a request and never responds and not for congestion control.

Dependent LUs

VTAM extends support for dependent (non-LU 6.2) LUs to APPN networks and combined subarea and APPN networks. Dependent LUs are owned by VTAM and rely on the SSCP to start sessions. There are SSCP-PU and SSCP-LU sessions for dependent LUs. Dependent LUs can start a session when the primary LU supports SLU-initiated sessions (for example, as does a VTAM application program). Dependent LUs can be the secondary LU in sessions with VTAM application programs in APPN, subarea, and combined APPN and subarea networks.

Note: If bisynchronous (BSC) 3270 sessions are required, traditional subarea paths must be retained for routing these sessions. These same subarea paths can be shared by APPN sessions using VR-based TG, permitting coexistence of APPN and BSC sessions.

Dependent LU server

The dependent LU server (DLUS) function facilitates conversion from a subarea environment to an APPN environment, allowing you to maintain central management of remote dependent LUs while benefiting from APPN throughout a network.

Note: The DLUR and DLUS architecture supports dependent secondary LUs (SLUs) only. The AS/400 SPLS function is not supported by DLUR and DLUS.

Two LU 6.2 sessions (one inbound, one outbound) are established between a DLUS and a dependent LU requester (DLUR), and these LU 6.2 sessions are collectively known as the CPSVRMGR pipe. SSCP-PU and SSCP-LU session flows use the CPSVRMGR pipe. An SSCP-PU session is established between a VTAM network node and the PU that owns the dependent LU, and an SSCP-LU session is established between VTAM and the dependent LU. Session initiation flows for the dependent LU are sent over the SSCP-LU session, and VTAM can use subarea or APPN flows to initiate a session with the PLU. BIND and session data are then routed directly between the PLU and the dependent LU.

For example, in [Figure 116 on page 406](#), VTAMB is a dependent LU server for LU4 and LU5. SSCP-PU sessions are established between VTAMB and PU4 and between VTAMB and PU5, and SSCP-LU sessions are established between VTAMB and the dependent LUs (LU4 and LU5). Session initiation flows for the dependent LUs are sent over the SSCP-LU sessions. VTAMB uses APPN flows to initiate a session with a PLU that VTAM can communicate with through an APPN network (for example, LU2 and LU3), and subarea

flows to initiate a session with a PLU reachable through a subarea network (for example, APPLA1). The BIND and session data are then routed directly between the PLU and the dependent LUs.




Figure 116. VTAM functioning as a dependent LU server

To implement the dependent LU server function, consider the following situations:

- NODETYPE=NN must be specified for a VTAM DLUS.
- For DLUR-initiated PU activation, no system definition is required. The dynamic switched definition facility can be used to dynamically define the PU. If the DLUR is adjacent, a PU name must be added to the REQACTPU by the DLUR. When the DLUR is adjacent, a CPNAME is sent in the REQACTPU and used to search for a PU. The previously defined CPNAME is found, and the REQACTPU is rejected. For information about the dynamic switched definition facility, see [Chapter 8, “Defining resources dynamically,”](#) on page 167 and [“Dynamic switched definitions”](#) on page 167.
- Dynamically defined PUs can be displayed using the DISPLAY ID command, specifying the DLUR CP name on the ID operand. This shows all PUs that currently have an SSCP-PU session (both dynamic and predefined PUs). You can also determine which LUs currently have an SSCP-LU session by displaying a specific PU.
- For VTAM to initiate activation of a PU, define the dependent LU requester by including the DLURNAME and DLCADDR operands on the PATH definition statement in the switched major node (regardless of whether its connection to the DLUR is switched or leased). DLURNAME specifies the CP name of the DLUR that owns the PU, and DLCADDR includes data link control (DLC) information used by the DLUR to locate the PU.
- In some cases the DLUS and the network node server of the DLUR are different VTAM nodes. In this case, there must be a searchable APPN path from the DLUS to the NNS of the DLUR. This path can go through multiple nodes if necessary.

Note: If the DLURNAME is not fully qualified, the NETID of the DLUS is used.

A sample switched major node to match Figure 116 on page 406 is shown in Figure 117 on page 407. DLUR1 is used in the switched major node as the name of the DLUR shown in Figure 116 on page 406.

- The DLUS and DLUR must be connected over an APPN path that supports the CPSVRMGR Class of Service.
- Takeover and giveback is supported for DLUR-attached PUs. The PU must be defined in VTAM with ANS=CONT, and the DLUR must be capable of handling ANS=CONT support. Giveback processing for DLUR-supported PUs occurs when an operator command is issued from the DLUS to relinquish control of the PU (for example, V NET,INACT,ID=dlur_pu,G), or an outage occurs for the CPSVRMGR pipe. Takeover processing can occur automatically for redial or can be initiated by the operator with a VARY DIAL command, for PUs previously released by giveback processing.

SAMPSWMN	VBUILD	TYPE=SWNET	
PU4	PU	ADDR=01,	X
		IDBLK=05D,	X
		IDNUM=00001,	X
		USSTAB=AUSSTAB,	X
		MODETAB=AMODETAB,	X
		ISTATUS=ACTIVE	X
PATHU4	PATH	PID=1,	X
		DLURNAME=DLUR1,	X
		DLCADDR=(1,C,INTPU),	X
		DLCADDR=(2,X,05D00001)	
LU4	LU	LOCADDR=1,	X
		PACING=(1,1),	X
		VPACING=2*	
PU5	PU	ADDR=02,	X
		IDBLK=05D,	X
		IDNUM=00002,	X
		USSTAB=AUSSTAB,	X
		ISTATUS=ACTIVE	X
PATHU5	PATH	PID=1,	X
		DLURNAME=DLUR1,	X
		DLCADDR=(1,C,INTPU),	X
		DLCADDR=(2,X,05D00002)	
LU5	LU	LOCADDR=1,	X
		PACING=(1,1),	X
		VPACING=2,	X
		MODETAB=AMODETAB	

Figure 117. Switched major node for a dependent LU server

Notes:

1. If using NCP, NCP Version 6 Release 2 or later is required.
2. The VARY INACT,TYPE=GIVEBACK command is used to release a DLUR-attached switched PU (and its subordinate LUs) from SSCP ownership. To use this function, the DLUR must support ANS=CONT.
3. A DLUS can serve multiple DLURs simultaneously.
4. A DLUR can be served by multiple DLUS VTAMs simultaneously.
5. Multiple DLURs can support the same PU (one at a time). The SSCP-PU session is through only one DLUR at any given time.
6. DLUS function cannot be used with XRF or cryptography in an XRF environment.
7. Redial occurs as follows for DLUS supported PUs:
 - When a DLUS initiates the activation of a PU but receives a negative response, VTAM attempts to redial over each valid path statement for the PU until successful or until all valid paths have been tried. Redial does not occur if the negative response indicates that the PU is already active, or that the fully-qualified PCID is not unique. If the fully-qualified PCID is not unique, the DLUS attempts to redial the PU over the same path with a newly generated PCID.
 - Redial does not occur if PU activation was initiated by the DLUR.
 - When a protocol violation, TDU error, or CPSVRMGR session outage signal is received for a particular DLUR, VTAM attempts to redial every active or pending-active PU served by that DLUR for which a valid PATH statement is found. If the PU is already active, VTAM performs the following actions:

- If the PU was defined with ANS=CONT and the DLUR supports this function, giveback processing is performed before attempting redial.
 - If the PU was defined with ANS=STOP or the DLUR does not support ANS=CONT, the PU is deactivated before attempting redial.
8. When a CP-CP session between a DLUS and DLUR fails, the CPSVRMGR session is deactivated, enabling reactivation of the CP-CP session.
 9. If the DLUS is located in a different APPN subnetwork from that of the DLUR or the PLU, APPN multiple network connectivity support is required (see [“Dependent LU server support across subnetwork boundaries”](#) on page 408). If CPSVRMGR pipe initiation is attempted across an APPN subnetwork boundary and the DLUR does not support this function, a protocol violation is detected and the pipe is terminated.
 10. Information regarding DLUR-attached resources can be processed by the configuration services XID exit. This exit routine can be coded so that VTAM either processes or denies requests for contact from known switched devices, or so that VTAM processes or denies requests for PU activation from a DLUR. For details, see [z/OS Communications Server: SNA Customization](#).
 11. The IDBLK/IDNUM specified on the PU definition statement must match the value specified on the DLCADDR operand (when DLCADDR information is provided for IDBLK/IDNUM) on the associated PATH statement.
 12. When a DLUR-attached dependent LU starts a session, or reports an active session as part of DLUR takeover, VTAM (as the DLUS) receives session information about the SESSST and passes that information to the network logical data manager (NLDM). However, if the session uses HPR and a path switch occurs, VTAM will not be aware of the change unless VTAM is at one end of the HPR route or if DLURSAW=YES is used. To supply NLDM with the updated path switch data between the DLUR and the other end of the session, even when VTAM is not at one end of the HPR route, use the DLURSAW=YES (this is the default) start option. The DLUR must be capable of reporting this information and be configured to do so.
 13. For information about the interaction between the DLURSAW and SAVERSCV start options, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Dependent LU server support across subnetwork boundaries

To use dependent LU server support across subnetwork boundaries, APPN multiple network connectivity support is required. NCP Version 7 Release 1 is also required for DLUS support across subnetwork boundaries.

If a DLUS is located in a different APPN subnetwork than a DLUR, the subnetwork of the DLUS must enable APPN multiple network connectivity support, as shown in [Figure 118 on page 409](#).

Note: The node on the border of the DLUR subnetwork (in [Figure 118 on page 409](#), the network node in subnetwork B) can be a network node, peripheral border node, or extended border node.

Also shown in [Figure 118 on page 409](#), if a DLUS is located in a different APPN subnetwork than the primary LU, both the subnetwork of the DLUS and the subnetwork of the primary LU must enable APPN multiple network connectivity support.

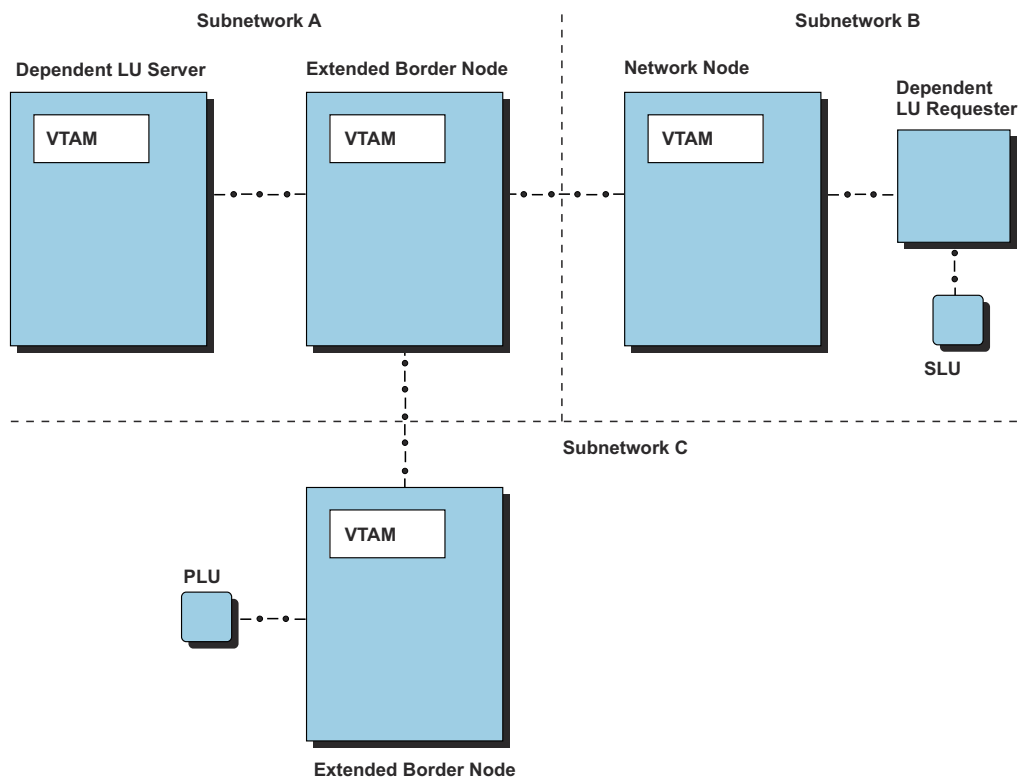


Figure 118. DLUS located in different APPN subnetwork than DLUR or PLU

If a primary LU exists in or through a subarea network, as shown in Figure 119 on page 410, all other APPN subnetworks through which the primary LU might be reached that use interchange nodes for connectivity to or through the subarea network must enable APPN multiple network connectivity support.

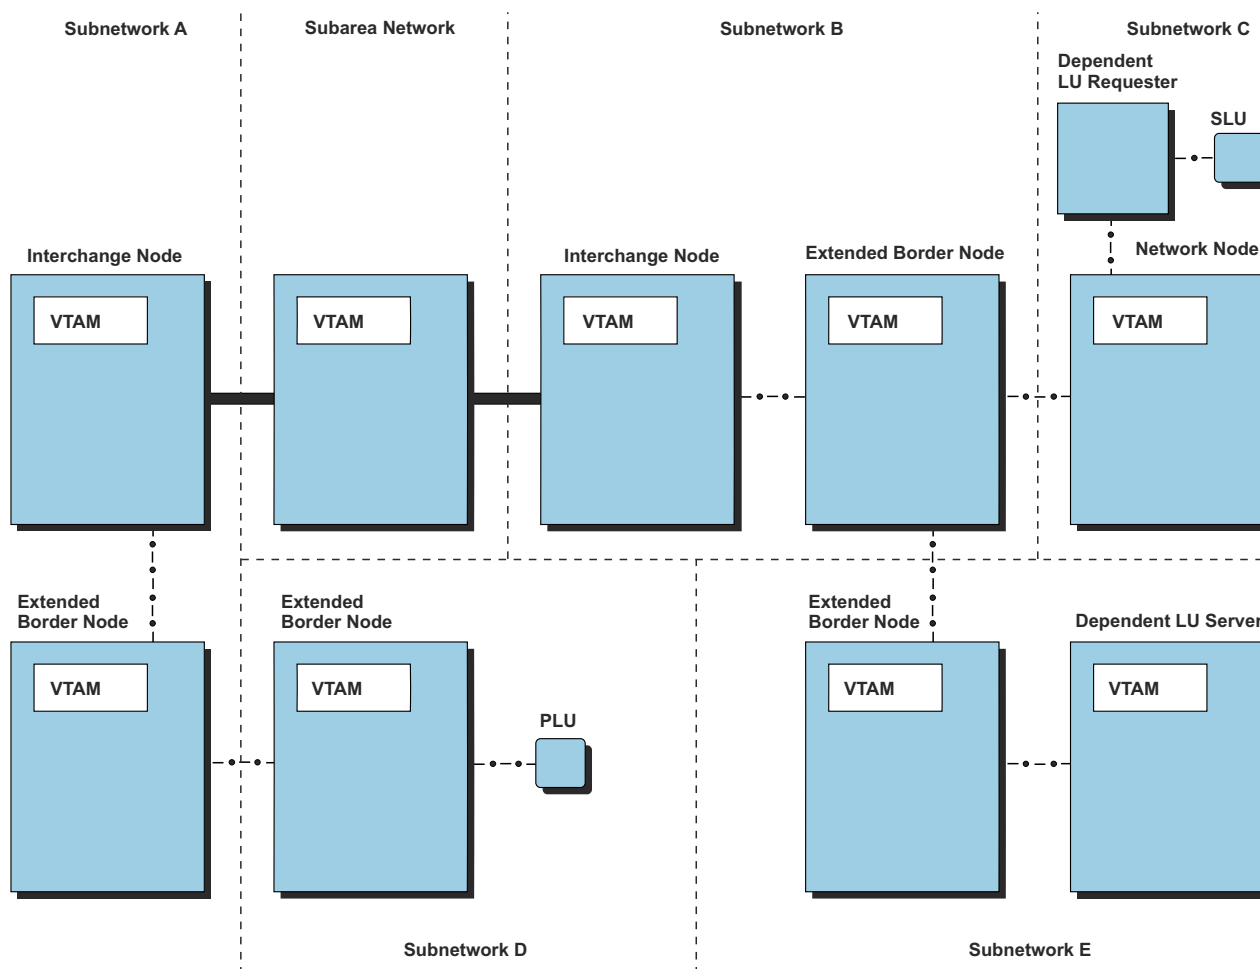


Figure 119. PLU exists in or through a subarea network

For further information about APPN multiple network connectivity support, see [“APPN multiple network connectivity”](#) on page 75.

Defining CDRSCs

When a resource is found in or through an APPN network, the interchange node through which the APPN network is entered becomes the owning CDRM of the resource to other VTAMs in the subarea network, unless the resource was predefined as a CDRSC with an owning CDRM name coded. If there are multiple interchange nodes, the owning CDRM can change, depending on the best search route at the time. If cross-domain resources (CDRSCs) are defined with the CDRM operand, they might have to be modified to specify an interchange node as the owning CDRM.

SSCP takeover

In a combined APPN and subarea network, VTAM provides full takeover capability.

A connection that is established as an APPN or LEN connection remains such until it is deactivated. However, an SSCP taking over an APPN connection views it as an LEN connection if the takeover SSCP has not implemented APPN. No CP-CP sessions can be established over the connection until the original SSCP regains control.

For example, in Figure 120 on page 411, if HOSTA fails and HOSTB takes over, the connection from the NCP to the adjacent CP is viewed by HOSTB as an LEN connection if HOSTB has not implemented APPN. No CP-CP sessions can be established over the connection until HOSTA regains control.




Figure 120. SSCP takeover when adjacent CP does not support CP name change

If the takeover SSCP has implemented APPN, a connection is still treated as an LEN connection if both of the following are true:

- The takeover SSCP is not the original SSCP, in which case the adjacent CP has a connection to a different SSCP after takeover (CP name change).
- The adjacent CP does not support the changing of CP names.

Although the adjacent CP still views the connection as an APPN connection with the original SSCP, the takeover SSCP views the connection as an LEN connection. Therefore, no CP-CP sessions can be established over the connection until the original SSCP regains control.

For example, in Figure 120 on page 411, if HOSTA fails and HOSTB takes over, the CP name (from the perspective of the adjacent CP) changes from HOSTA to HOSTB. If the adjacent CP does not support the changing of CP names, HOSTB treats the connection to the adjacent CP as an LEN connection (even if HOSTB has implemented APPN) and no CP-CP sessions can be established over the connection until HOSTA regains control. If the adjacent CP does support the changing of CP names, CP-CP sessions can be established when HOSTB takes over.

If the adjacent CP is located over a subnetwork boundary, the following conditions must also be met to allow a new CP-CP session with the takeover SSCP:

- The takeover SSCP must have APPN multiple network connectivity support.
- The definitions in the takeover SSCP must be consistent with those in the original SSCP (for example, a subnetwork boundary cannot be changed to a native connection during takeover).

If these conditions are not met, the takeover SSCP views the connection as an LEN connection. For information about APPN multiple network connectivity support, see [“APPN multiple network connectivity”](#) on page 75.

If the adjacent control point is another composite network node as shown in [Figure 121 on page 412](#), SSCP takeover with APPN connectivity occurs if the takeover SSCP has implemented APPN. For example, in [Figure 121 on page 412](#), if HOSTB has implemented APPN and HOSTA fails, HOSTB can take over, and a CP-CP session between the composite network nodes (HOSTC/NCP2 and HOSTB/NCP1) is established. The APPN connection between the NCPs is viewed as an LEN connection by the adjacent control point after takeover if the takeover SSCP has not implemented APPN.




Figure 121. SSCP takeover when adjacent CP is another composite network node

Notes:

1. In [Figure 121 on page 412](#), if HOSTA takes over for HOSTC or HOSTC takes over for HOSTA, APPN and LEN connectivity is not available after takeover (assuming a CP-CP session originally existed between only the composite network nodes shown) because the takeover SSCP cannot have a connection with itself (that is, after takeover, the takeover SSCP and the NCPs now form a composite network node). APPN and LEN connectivity is not available until the failing host regains control or another host (for example, HOSTB) takes control of the NCP originally owned by the failing host.
2. In [Figure 121 on page 412](#), a subarea connection and an APPN connection exist between the NCPs. If both the subarea connection and the APPN connection are active concurrently, the route used for communication between the NCPs is unpredictable. To control which link is used, you can choose to activate only one connection. For SSCP takeover, however, the subarea connection must be activated if not already active.

Establishing and controlling sessions

In a combined APPN and subarea environment, an independent LU can be accessible over APPN and LEN connections. In this case, VTAM always gives preference to APPN connections, if the connections

are known and available. APPN connections are known when an APPN-capable PU name or ISTAPNPU is included in the ALSLIST for the resource. If no APPN connection is available, an LEN connection is used (if available).

To cause VTAM to sometimes use an LEN connection or sometimes use an APPN connection, use the adjacent link station selection function of the session management exit routine. For information about coding this exit routine, see *z/OS Communications Server: SNA Customization*.

To force VTAM to use particular LEN connections and not allow a search of the APPN network, include only those LEN adjacent link stations on the ALSLIST operand and also code ALSREQ=YES for the resource. Session setup will continue using an LEN adjacent link station from the list even if a route exists through the APPN network.

Controlling searches

When trying to locate a DLU, an interchange node or migration data host can search the APPN network, the subarea network, or both. The order that VTAM searches is based on the defined ADJSSCP tables, start option values, and the origin of the search request. Using the ADJSSCP tables and start options, you can concentrate the search on the part of the network containing the most frequently sought resources. See [“Adjacent SSCPs” on page 431](#) for information about using ADJSSCP tables.

For further information about searching an APPN network, see [“Network routing and resource location for APPN nodes” on page 235](#).

For further information about searching a subarea network, review the appropriate sections of Chapter 17, [“Implementing a subarea network,” on page 421](#), beginning with [“Start options defining other domains” on page 422](#).

Using SORDER to control network search order

SORDER controls the order in which the APPN and subarea portions of the network are searched when adjacent SSCP routing is performed for a search request received by an interchange node or migration data host. SORDER can be specified as a start option or coded on individual ADJSSCP table definitions (or both). In general, the SORDER start option value should be chosen based on where the most frequently sought after resources are located (in the APPN network or in the subarea network). Although this approach provides the best search performance for MOST resources in the network, a single SORDER value rarely results in optimal search performance for ALL resources. For these cases, the search order can be specified on a more granular level by coding SORDER individually on each ADJSSCP table.

The default value for the SORDER start option (SORDER=APPN) will cause VTAM to prefer searching the APPN network before searching the subarea network (adjacent CDRMs). However, if the CDRM that owns the target resource is known (defined on the cross-domain resource definition or learned from previous searches), new searches will be sent to this CDRM first, before searching the APPN network. If the APPN network should always be searched first even before the owning CDRM (for example, to maximize the use of High Performance Routing), SORDER=APPNFRST should be used.

If most target resources are still located in the subarea portion of the network, SORDER=SUBAREA should be used. In this case, the APPN network will be searched only after all adjacent CDRMs have been searched. If VTAM has been enabled for APPN but there are currently no resources located in the APPN network (that is, no APPN searching should be performed), SORDER=ADJSSCP should be used. With SORDER=ADJSSCP, the APPN network is only searched if an ISTAPNCP ADJCDRM entry has been explicitly coded in the ADJSSCP table that is chosen for the search.

To specify the search order for individual adjacent SSCP tables, code the SORDER operand on the NETWORK or CDRM statements that define each adjacent SSCP table. If SORDER is coded on a NETWORK statement, that value will sift down to all subordinate adjacent SSCP tables until the next NETWORK (or ADJLIST) statement is encountered. To override the SORDER value on the preceding NETWORK statement, code the SORDER operand on the CDRM statement that defines the subordinate adjacent SSCP table. To override the SORDER value on the preceding NETWORK statement with the current SORDER start option value, code SORDER=STARTOPT on the CDRM statement. (SORDER=STARTOPT is the default value for NETWORK statements.)

Notes:

1. If a search request is received by VTAM from the APPN network, SORDER is ignored. The search order used for searches received from the APPN network will depend on what information (if any) exists in the Directory Services database. The DISPLAY DIRECTORY command can be used to display information in the Directory Services database.
2. If SSCPORD=PRIORITY is specified and previous searches for the target resource have been performed, the search order can be modified based on the results of these previous searches. Use the DISPLAY ADJSSCP command with the CDRSC operand to display the search order that will be used by VTAM on the next search for this resource.
3. If SORDER=APPNFRST, APPN or SUBAREA is specified for an adjacent SSCP table, any ISTAPNCP ADJCDRM entries coded within that table are ignored, because VTAM will automatically place the ISTAPNCP ADJCDRM entry in the appropriate position in the table. This is true regardless of where the SORDER value was obtained.
4. If SORDER=ADJSSCP is specified for an adjacent SSCP table and an ISTAPNCP ADJCDRM entry is not coded within that table, the APPN network will not be searched (unless SSCPORD=PRIORITY is specified and a previous search for this resource determined that it could be found in the APPN network).

Using SSEARCH to limit subarea network searches

The SSEARCH start option controls whether the subarea network is searched when a network search request is received from the APPN network. The default value (SSEARCH=YES) means that subarea network searching is allowed. In cases where both the APPN network and the subarea network provide connectivity to (or toward) the target resource, use SSEARCH=APPNFRST to force the APPN network to be used whenever possible (for example, to maximize the use of High-Performance Routing). SSEARCH=APPNFRST works like SSEARCH=YES, except that the subarea network is only searched after all APPN searching has been performed.

If APPN searches should only be propagated into the subarea network when the target resource has been defined (as a cross-domain resource) or has been previously found in the subarea network, SSEARCH=CACHE should be specified. To prevent all APPN searches received by this VTAM from entering the subarea network, use SSEARCH=NO.

Using the CDRSCTI start option to reduce broadcast searches of APPN

The setting of the CDRSCTI start option controls the amount of time that a dynamic CDRSC is retained by VTAM after its last session ends. If a dynamic CDRSC has been freed by VTAM in a combined environment, search requests originating in the subarea network or in the domain of the interchange node itself can result in unnecessary broadcasts to locate a subarea resource in the APPN network.

Therefore, to reduce broadcast searches of an APPN network that does not have a central directory server, you can increase the CDRSCTI value from the default of eight minutes to as much as seven days at each interchange node. This will minimize the potential for unnecessary broadcasts of the APPN network, because dynamic CDRSCs are retained for a longer period of time after the last session ends and VTAM has more of an opportunity to take advantage of priority routing to locate the resource in the subarea network.

Using the DISJOINT operand to define disjoint subarea networks

The DISJOINT operand on the CDRM definition statement in the cross-domain resource manager major node can also be used to control session request searches. At an interchange node, DISJOINT indicates whether an adjacent CDRM is disjoint from other SSCPs with the same NETID as the adjacent CDRM, and the only connection between the disjoint SSCPs is through the APPN network.

For example, in [Figure 122 on page 415](#), HOSTA and HOSTB both have the same network identifier (NETA), but their only connection to each other is through adjacent network NETB. When a session request is received by HOSTC from HOSTA, HOSTC might decide to search the attached APPN network. When the search is forwarded into the APPN network, HOSTC includes information about the adjacent

SSCP from which the request was received; namely, the HOSTA network identifier (NETA) and an indication of whether HOSTA is disjoint from other SSCPs within NETA (as defined by HOSTC using the DISJOINT operand in the CDRM definition statement for HOSTA).




Figure 122. Disjoint hosts

When the APPN search request is received by HOSTD, HOSTD uses this information to determine the SSCPs to which the request should be routed. HOSTD routes the request to other SSCPs that are not in NETA (the network identifier provided by HOSTC, the interchange node serving as the APPN entry point).

However, when deciding whether to route the request to SSCPs within NETA (the network identifier provided by HOSTC), HOSTD uses the DISJOINT network indicator provided by HOSTC to make the decision. If DISJOINT=YES was coded in HOSTC (and provided to HOSTD on the APPN search request), HOSTD routes the search request to other SSCPs within NETA, because it is assumed that the APPN network can provide the only connectivity between the two disjoint parts of network NETA. If DISJOINT=NO was coded (or allowed to default) in HOSTC, HOSTD does not route the request to other SSCPs within NETA, because it is assumed that other connectivity exists between HOSTA and HOSTB (see Figure 123 on page 416). When HOSTD and HOSTC fail to locate the target LU, HOSTA continues to try other SSCPs and find the direct path to HOSTB.

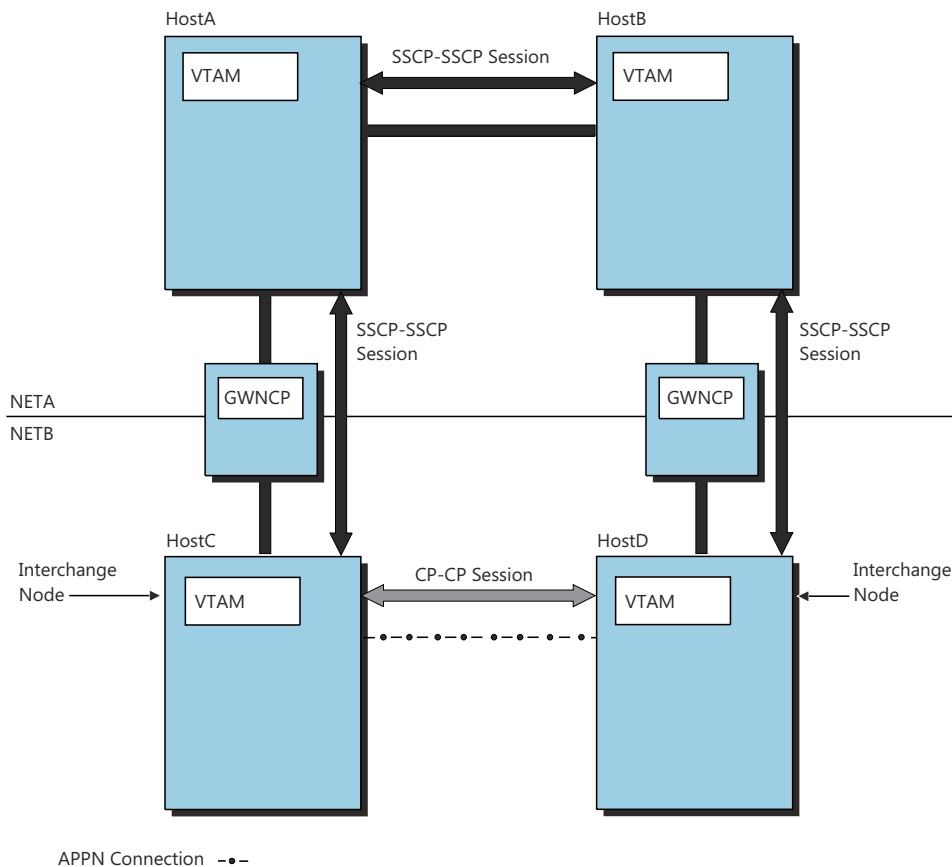


Figure 123. Hosts with subarea connection

Note: The preceding information about the DISJOINT operand also applies when the net IDs of all four hosts in Figure 123 on page 416 and Figure 124 on page 416 are the same (NETA).

APPN and subarea Class of Service resolution

In a subarea network, the LOGMODE resolves to a subarea COS at the SLU SSCP. In Figure 124 on page 416, the following cases occur in a subarea environment:

- For PLU-initiated sessions, the COS is resolved at VTAM2 when the CDINIT is received from VTAM1 before returning the CDINIT response.
- For SLU-initiated sessions, the COS is resolved at VTAM2 before the CDINIT is sent to VTAM1.



Figure 124. LOGMODE resolution example

In an APPN network, the LOGMODE is always resolved to an APPN COS at the CP(OLU) or NNS(OLU). Using Figure 124 on page 416, the following cases occur in an APPN environment:

- For PLU-initiated sessions, the APPN COS is resolved at VTAM1.

- For SLU-initiated sessions, the APPN COS is resolved at VTAM2.

This difference between the architectures (subarea and APPN) creates new considerations in the implementation of mixed APPN and subarea networks, especially when you consider the fact that VTAM allows for different LOGMODE tables to be defined for various LUs. For example, consider a PLU-initiated session in both a subarea and an APPN network configuration.

In a subarea network

A CDINIT flows from VTAM1 to VTAM2. VTAM2 (the owner of the SLU) uses the SLU definition to resolve the LOGMODE to a subarea COS. The SLU definition can specify a specific LOGMODE table.

In an APPN network

The APPN COS must be determined before sending an APPN Locate from VTAM1 to VTAM2. VTAM1 uses the LOGMODE table associated with the SLU to resolve the LOGMODE to an APPN COS. The SLU is defined to VTAM1 as a CDRSC. If the CDRSC has been statically defined, the logon mode table to be used can be identified on the MODETAB operand. If the CDRSC has been dynamically created, the default logon mode table (ISTINCLM) is used.

If every possible SLU is predefined as a CDRSC in every VTAM that has a PLU that might wish to start a session with these SLUs, the MODETAB operand can be added to these CDRSCs to make sure the correct LOGMODE table (and, therefore, APPNCOS) is selected. However, requiring SLU definitions at every possible PLU VTAM is both undesirable and unreasonable in an APPN environment. Using dynamic CDRSCs for these SLUs is easier to maintain. Without any definition changes, the default logon mode table (ISTINCLM) is used for dynamic CDRSCs.

In addition, to ensure required logmode to COS resolution, all SLUs need to be defined at any VTAM that is an intermediate network node on the BIND path. For example, if a composite network node existed between VTAM1 and VTAM2, the CNN would have to be able to resolve the LOGMODE name to a subarea COS so that an appropriate ER and VR can be chosen through the CNN. Because neither the PLU nor SLU has been predefined to the CNN, dynamic CDRSCs are created for both and the SLU would inherit ISTINCLM as its logon mode table.

To ensure that the required COS is being used when LU-LU session requests use both APPN links and subarea VRs, there are several choices available for specifying the appropriate APPN and subarea COS for a given logon mode name. The following choices are provided in order of the preference used by VTAM for COS selection.

1. You can define APPN-to-subarea (APPNTOSA) and subarea-to-APPN (SATOAPPN) Class of Service mapping tables to define the mappings between APPN and subarea Classes of Service.
2. You can modify your customer-defined logon mode tables to specify the appropriate APPN and subarea Class of Service names for each logon mode table entry. (You can also associate one of your customer-defined logon mode tables with all dynamic CDRSCs using the DYNMODTB start option.)
3. You can modify the IBM-supplied default logon mode table, ISTINCLM, to specify (or change) the APPN and subarea Class of Service names for the IBM-supplied logon modes. (This might be required if you allow VTAM to create dynamic CDRSCs and you have not specified a logon mode table name on the DYNMODTB start option.)

Resolving logon mode names to subarea and APPN Classes of Service

VTAM uses the logon mode table for the secondary LU (SLU) to initially resolve the logon mode name to a subarea COS. This happens at the VTAM that owns the SLU or at the Interchange Nodes (ICN) that must send a session request into the APPN network.

For sessions that will use APPN links, one of two methods is used to determine the APPN COS for the session. If a subarea-to-APPN (SATOAPPN) Class of Service mapping table has been defined at the VTAM that must choose the APPN Class of Service, it will be used to map the resolved subarea COS to an appropriate APPN COS. If a SATOAPPN Class of Service mapping table has not been defined, or a table is found with no matching entry, VTAM uses the logon mode table for the SLU (as defined in that VTAM node) to resolve the logon mode name to an appropriate APPN COS.

The same is true for session requests (or replies) received by VTAM from an APPN network that will also use subarea virtual routes. One of two methods is used to determine the appropriate subarea COS for the

session. If an APPN-to-subarea (APPNTOSA) Class of Service mapping table is defined at the VTAM that must choose the subarea COS, it will be used to map the APPN COS received on the APPN session request (or reply) to an appropriate subarea COS. If an APPNTOSA Class of Service mapping table is not defined, or a table is found with no matching entry, VTAM uses the logon mode table for the SLU (as defined in that VTAM node) to resolve the logon mode name to an appropriate subarea COS.

Note: VTAM always uses the logon mode table for the secondary LU (SLU) to resolve the logon mode name to an APPN or subarea Class of Service, when needed. But at interchange nodes and for other node roles, the SLU might be a dynamic CDRSC. In this case, the IBM-supplied default logon mode table (ISTINCLM) would normally be used. If you want to have a different table used for these resources, you can use the DYNMODTB start option to specify a different logon mode table to be used for all dynamically created CDRSCs.

Resolving the logon mode name to session parameters and subarea COS

Session parameters and subarea Class of Service are determined in the SSCP that is closest to the secondary logical unit (SLU). This includes cases where this VTAM:

- Owns the SLU (the SLU is an application or a dependent LU in this VTAM domain), like HostD owns LUD in [Figure 125 on page 419](#)
- Owns the boundary function in the direction of an LEN-attached independent SLU, like HostD does for LUD in [Figure 125 on page 419](#)
- Owns the boundary function in the direction of an APPN-attached independent SLU, like HostA does for LUD in [Figure 125 on page 419](#). This VTAM might:
 - Own the primary logical unit (PLU), and the SLU is elsewhere in (or through) the APPN network.
 - Be an interchange node, with the PLU in (or through) the subarea network, and the SLU elsewhere in (or through) the APPN network.
 - Be part of a composite network node that is an intermediate node on the BIND path of the session. In this case, the Class of Service is needed to determine the subarea route through the composite network node.

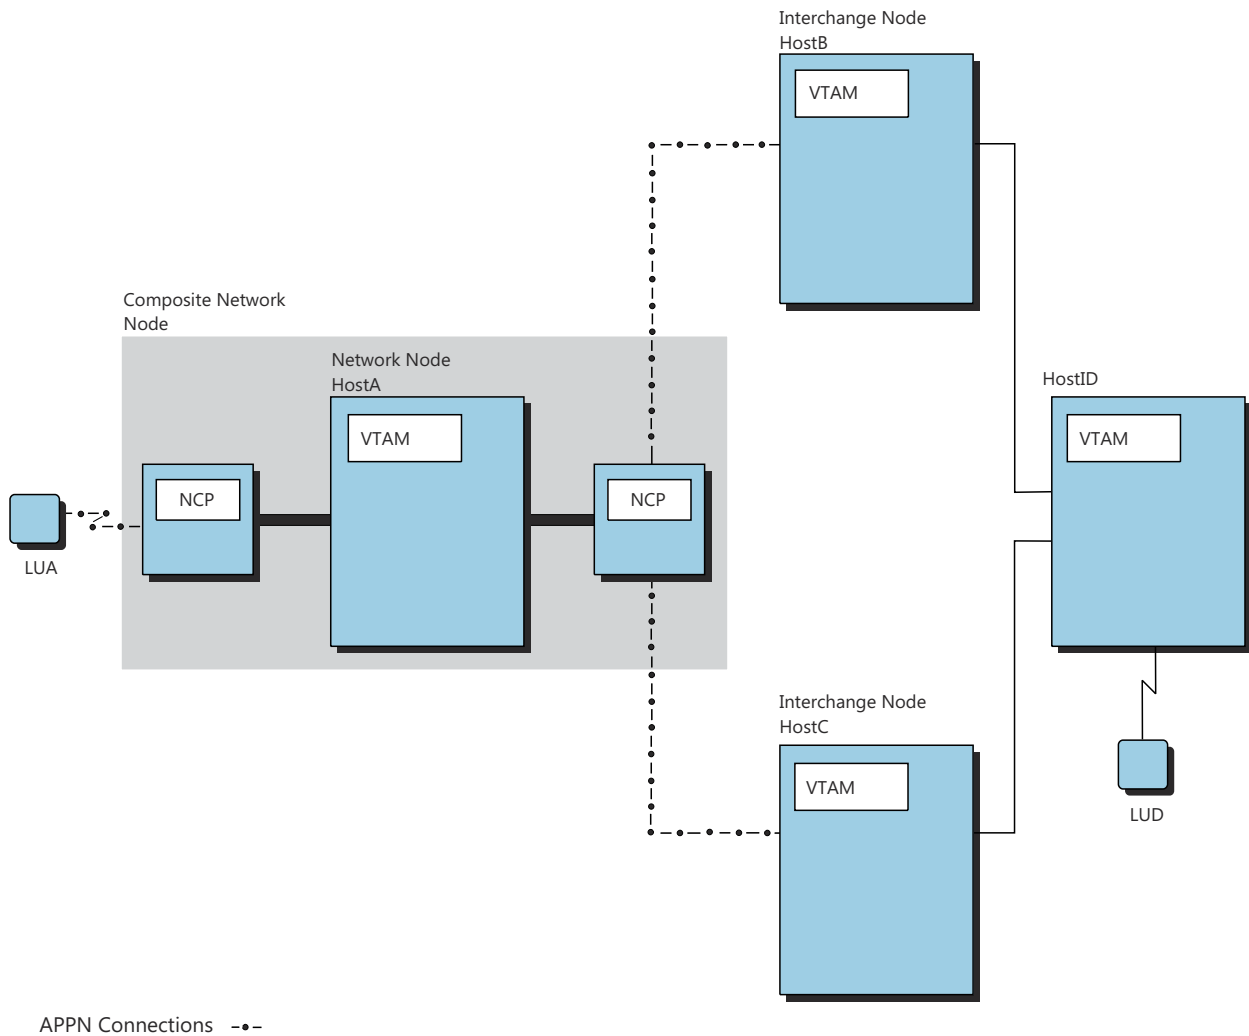


Figure 125. Class of Service resolution at multiple nodes

Resolving the logon mode name to an APPN Class of Service

The APPN Class of Service is determined in the APPN control point (CP), closest to the originating logical unit (OLU) for the session, that has the ability to perform the determination. For VTAM, this includes cases where this VTAM:

- Has the OLU in its domain and the destination logical unit (DLU) is elsewhere in (or through) the APPN network, like HostA in [Figure 125 on page 419](#)
- Is the network node server for an end node (EN) that owns the OLU but is not capable of determining the APPN Class of Service itself, and the DLU is owned by this network node server or is elsewhere in (or through) the APPN network (like HostA in [Figure 125 on page 419](#))
- Is an interchange node, with the OLU in (or through) the subarea network and the DLU elsewhere in (or through) the APPN network, like HostB or HostC in [Figure 125 on page 419](#)

Defining APPN and subarea Classes of Service in logon mode tables

For each logon mode table associated with LUs that might be involved in a session with both APPN links and subarea virtual routes, specify an APPN Class of Service by adding the APPNCOS operand and a subarea Class of Service using the COS operand on the appropriate MODEENT macroinstructions. The subarea virtual routes use the Class of Service name coded on the COS operand, and APPN links use the Class of Service name coded on the APPNCOS operand. Using both the COS and APPNCOS operands

enables you to use two different Class of Service names to represent the same Class of Service within the two subnetworks. This can be very useful as you migrate to APPN.

If neither the COS operand nor the APPNCOS operand is coded, the default Class of Service names (the unnamed Class of Service entry for subarea and the #CONNECT entry for APPN) are used.

If the COS operand is coded and the APPNCOS operand is not coded, the Class of Service name specified on the COS operand is used for both the subarea and APPN sessions. Therefore, when you want to use the same Class of Service name to represent identical Classes of Service in both the APPN and subarea networks, you do not need to code the APPNCOS operand when the COS operand is already coded. However, make sure that a corresponding APPN Class of Service definition exists in all network nodes. It is recommended that the architected standard APPN COS names be used.

If you code the APPNCOS operand and the COS operand is not coded, the unnamed Class of Service is used for the subarea session, and the Class of Service designated on the APPNCOS operand is used for the APPN session.

Defining APPNTOSA and SATOAPPN class of service mapping tables

Rather than resolving the logon mode name to a subarea or APPN class of service at each APPN or subarea network boundary, you can choose to map subarea classes of service to APPN classes of service, and vice versa. This is done by coding a subarea-to-APPN COS mapping table (SATOAPPN) and an APPN-to-subarea (APPNTOSA) COS mapping table. These tables are used at any VTAM that must choose an APPN or subarea Class of Service when the corresponding Class of Service (subarea or APPN) is already known.

Note: If subarea-to-APPN or APPN-to-subarea COS mapping tables are defined, these tables are used instead of resolving the logon mode name to an appropriate APPN or subarea COS using the logon mode table for the secondary LU. If COS mapping tables are not defined, or a matching entry is not found, the logon mode table for the SLU is used to resolve the logon mode name to an appropriate APPN or subarea Class of Service.

Adding and moving nodes

To add or move workstations or applications in an APPN network, no coordinated LU definitions are required. As workstations and applications first become available or are later moved, APPN directory functions can locate them dynamically. If new hosts are added, only the resources on that host must be defined to that host.

If an APPN network node is to be moved from one APPN subnetwork to another, purge its topology database before activating CP-CP sessions in the new APPN subnetwork. This prevents broadcasting of network topology from the original subnetwork to the new subnetwork, information that is of no use in the new subnetwork. For a description of APPN subnetworks, see [“APPN multiple network connectivity” on page 75](#).

No definitions are necessary in the subarea network for resources added or moved in the APPN network. APPN directory functions can locate them dynamically.

Chapter 17. Implementing a subarea network

This topic contains information needed when more than one VTAM subarea domain exists in the network. Chapter 3, “Implementing a VTAM network,” on page 19 contained basic concepts and issues requiring consideration for each VTAM domain. This topic expands on Chapter 3, “Implementing a VTAM network,” on page 19 to include multiple domain issues.

A multiple-domain subarea network contains more than one system services control point (SSCP). Control of the resources in the network is divided among the SSCPs. Defining your VTAM network involves identifying your domains and defining the resources in those domains to the SSCPs.

Note: A VTAM host contains a host CP, even if you have defined your VTAM as a subarea node and not as an APPN node. This host CP is used for management services transport and is not used for session setup or routing.

Figure 126 on page 421 is the same as the configuration in Figure 4 on page 20, except that this is a multiple-domain network because it includes another host. Note that NCP12 is now NCP21. This is for naming convention purposes.




Figure 126. Multiple-domain network

Code everything that you would code to define the network in Figure 4 on page 20, except that you need to put the definitions for NCP21 and its peripheral nodes in HOST2 rather than in HOST1. You also need to code at least the following in HOST2:

- Start options describing HOST2 (NETID, HOSTSA, SSCPID, and SSCPNAME).
- One application program minor node defining APPL2.
- One NCP major node defining NCP21. This includes the NCP definitions. In NCP21 major node, also code another PATH definition statement to attach to subarea 2, and change the PATH definition statement for NCP21 in NCP11 major node to DESTSA=21.

- VTAM PATH definition statements defining the explicit and virtual routes used for data flow throughout the network.
- One cross-domain resource manager major node containing two CDRM statements (one for HOST1 and one for HOST2).
- One cross-domain resource major node (or indicate the use of dynamic definition of cross-domain resources in the CDRM major node).
- One adjacent SSCP table (or indicate dynamic definition of adjacent SSCPs using start options [SSCPDYN and DYNASSCP]).

Also code at least the following in HOST1:

- The HOSTSA start option
- A cross-domain resource manager major node containing two CDRM statements (one for HOST1 and one for HOST2)
- A cross-domain resource major node (or indicate use of dynamic definition of cross-domain resources in the CDRM major node)
- An adjacent SSCP table (or indicate dynamic definition of adjacent SSCPs using start options [SSCPDYN and DYNASSCP])

Start options defining other domains

In a multiple-domain subarea network, you can define the location of cross-domain resource managers and also let VTAM define cross-domain resource managers dynamically.

Defining the location of cross-domain resource managers by coding adjacent SSCP tables

In a multiple-domain network, session setup involves VTAM locating resources (session partners) in the network. Resources are located by routing a session initiation request to VTAMs in other domains. You can define a list of cross-domain resource managers (CDRMs) to which a session setup request can be routed. The list of CDRMs that you code for the search (referred to as adjacent SSCPs) can be supplemented dynamically by your VTAM domain. The search order, and the number of VTAMs to which the session request is routed, can affect session setup time. For more information, see [“Adjacent SSCPs” on page 431](#).

VTAM provides two start options, SSCP DYN and SSCP ORD, that control how VTAM processes adjacent SSCP tables. SSCP ORD can be specified as a start option or coded on individual ADJSSCP table definitions (or both).

SSCP DYN start option

SSCP DYN determines whether VTAM should add entries dynamically to the adjacent SSCP table (that is, supplement the list of adjacent SSCPs that you have coded with other SSCPs that it discovers dynamically). The addition of these entries to the list can increase the time needed to search for a cross-domain resource.

If you code SSCP DYN=NO, you might have to define additional adjacent SSCP tables. When an application program passes a session (issues a CLSDST PASS macro instruction) during session setup, VTAM does not build an adjacent SSCP table in the application program network. When SSCP DYN=YES is specified in the application program network, VTAM builds an entry in the table automatically.

SSCP ORD start option

SSCP ORD determines the order in which VTAM processes the entries in the table to route session requests to adjacent SSCPs. SSCP ORD can be specified as a start option, modified using MODIFY VTAM OPTS, or coded separately for each ADJSSCP table (or all three). If SSCP ORD is not coded for a given ADJSSCP table, the current value of the SSCP ORD start option value is used each time that ADJSSCP table is selected for use. SSCP ORD indicates whether VTAM is to process the list in the order in which it is defined (DEFINED) or whether it is to route the session request to the last successful entry (PRIORITY). VTAM remembers successful and unsuccessful adjacent SSCP entries for routing

session initiation requests. If you specify SSCPORD=PRIORITY, VTAM always uses the last successful adjacent SSCP entry to route the session setup request.

One exception to order of search, regardless of the SSCPORD value, is that VTAM always routes the session-initiation request directly to the VTAM that owns the resource if it has an SSCP-SSCP session with that owner.

Defining the location of cross-domain resource managers dynamically

If you use the start option DYNASSCP=YES, you do not have to code an adjacent SSCP table. VTAM dynamically builds the list in the order in which the SSCP-SSCP sessions are activated. VTAM then uses this dynamically constructed list to route session initiation requests to the CDRMs, if an appropriate list is not defined.

Specifying timeout values for locating cross-domain resources

The following start options also affect VTAM processing in a multiple-domain subarea network:

CDRSCTI start option

In a multiple-domain environment, you define resources that are controlled by a VTAM in another domain using a cross-domain resource (CDRSC) definition statement. However, if you do not code the CDRSC definition statement, you can specify that a CDRSC definition be dynamically created by using the CDRSC=OPT operand on the cross-domain resource manager (CDRM) definition statement that defines that VTAM. If your VTAM dynamically defines a cross-domain resource, that dynamic definition can be retained by VTAM for a user-specified period of time or deleted immediately after the last session for that cross-domain resource terminates. The cross-domain resource timeout value (CDRSCTI) start option determines the processing VTAM uses for dynamic cross-domain resources. The value of this start option specifies the amount of time that VTAM is to retain the control blocks for a dynamic cross-domain resource after all sessions with that resource are terminated. For more information, see [“Dynamic definition of cross-domain resources” on page 425](#).

IOPURGE start option

During cross-domain session initiation, a CDINIT is sent to the adjacent node that is most likely to be the resource owner. If a positive response is received, session setup continues with that adjacent SSCP. If a negative response is received indicating the resource was not found, the next adjacent SSCP that potentially owns the resource is obtained from the adjacent SSCP table and a CDINIT is sent to it. This continues until the resource is either found or the adjacent SSCP table is exhausted.

A situation could arise where the adjacent SSCP that was queried has a critical failure and cannot respond. In this case, the node that is trying to initiate the session waits forever for a reply (IOPURGE=0, the default), possibly blocking any future attempt to find the resource. Because the failing SSCP may not be the resource owner, the resource may be available and could be found if the search was allowed to continue through the adjacent SSCP table.

The IOPURGE start option can be used to specify how long VTAM will wait for a response before continuing a search. After a specified time interval, if a response has not been received, VTAM continues its search through the adjacent SSCP table until it either finds the resource or the table has been exhausted.

Identifying VTAMs in other domains (CDRMs)

A cross-domain resource manager (CDRM) is the part of an SSCP that supports cross-domain session setup and takedown. Before logical units in one domain can have cross-domain sessions with logical units in another domain, an SSCP-SSCP session must be established between the SSCPs of the two domains.

For an SSCP-SSCP session to exist, VTAM must know about all cross-domain resource managers with which it will communicate. You must define to VTAM its own cross-domain resource manager and all other cross-domain resource managers in the network. The cross-domain resource manager that represents the SSCP in your domain is called the host cross-domain resource manager. The cross-domain resource managers that represent the SSCPs in other domains are called external cross-domain resource managers.

Thus, to have a session between two SSCPs, define two cross-domain resource managers to each VTAM: one for the host and one for the external cross-domain resource manager. You file these definitions in a CDRM major node. Each cross-domain resource manager is a minor node.

One or more major nodes can be used to define cross domain resource managers. Each CDRM major node is defined with a VBUILD definition statement, and each minor node is defined with a CDRM definition statement.

Figure 127 on page 424 shows two types of major and minor nodes in a multiple-domain network.

Cross-Domain Resource Manager Major Node

Each set of cross-domain resource managers is a major node. In this example A01M is this host's CDRM. A02M and A03M represent host processors in other domains.

Each cross domain resource manager is a minor node.

Cross-Domain Resource Major Node

Each set of cross-domain resources is a major node.

Each cross-domain resource is a minor node. The definition for a cross-domain resource may optionally specify the owning CDRM.

VTAM (In Host Processor)




Figure 127. Major and minor nodes in multiple-domain environment

Following is an example of major node CDRM for host A01M. The same definition set can be used at A02M and A03M. In this case, the definition set is at A01M. Thus, A01M is the host cross-domain resource manager and A02M and A03M are external cross-domain resource managers.

A01M	VBUILD TYPE=CDRM CDRM CDRDYN=YES, CDRSC=OPT, ELEMENT=1, ISTATUS=ACTIVE, SUBAREA=1, VPACING=63	AUTHORIZE DYNAMIC CDRSC ALLOW DYNAMIC CDRSCs CDRM ELEMENT NUMBER VTAM INITIAL STATUS NETWORK UNIQUE SUBAREA ADDRESS PACING BETWEEN CDRMS
* A02M	CDRM CDRDYN=YES, CDRSC=OPT, ELEMENT=1, ISTATUS=ACTIVE, SUBAREA=2, VPACING=63	AUTHORIZE DYNAMIC CDRSC ALLOW DYNAMIC CDRSCs CDRM ELEMENT ADDRESS VTAM INITIAL STATUS NETWORK UNIQUE SUBAREA ADDRESS PACING BETWEEN CDRMS
* A03M	CDRM CDRDYN=NO, CDRSC=REQ, ELEMENT=1, ISTATUS=ACTIVE, SUBAREA=3, VPACING=63	DYNAMIC CDRSC NOT AUTHORIZED REQUIRE PREDEFINED CDRSCs CDRM ELEMENT ADDRESS VTAM INITIAL STATUS NETWORK UNIQUE SUBAREA ADDRESS PACING BETWEEN CDRMS

Each CDRM minor node definition includes its network address and its initial status (active by default, or inactive). CDRM minor nodes are assigned to the subarea address of the appropriate host.

You can use the same set of CDRM definition statements for the entire network. Each VTAM uses the subarea specified on the CDRM definition statements to determine which definition statement defines its

own host cross-domain resource manager. If the subarea on a CDRM definition statement matches the host subarea, VTAM uses the operands on this CDRM definition statement that apply to the host cross-domain resource managers and ignores operands that apply to external cross-domain resource managers. If the subarea on a CDRM definition statement does not match its host subarea, VTAM assumes that this CDRM definition statement applies to an external cross-domain resource manager. VTAM then uses only the operands that apply to external cross-domain resource managers.

When defining the host cross-domain resource manager, you can specify whether the host cross-domain resource manager is authorized to dynamically define cross-domain resources. If the host cross-domain resource manager and an external cross-domain resource manager are authorized in a VTAM domain, the other-domain SSCP can request a session setup for a logical unit for which no cross-domain resource (CDRSC) was defined. The receiving SSCP handles the request by dynamically creating a temporary CDRSC definition for the logical unit. Using dynamic definition of cross-domain resources can greatly reduce the number of definitions required in a multiple-domain network.

The default for the CDRDYN start option is YES, enabling dynamic definition of cross-domain resources. The CDRDYN operand can be coded on the host CDRM definition statement, but the CDRDYN start option overrides it. Also code CDRSC=OPT on external CDRM definition statements whose resources you want VTAM to dynamically define. The CDRSC operand is meaningful only on CDRM definition statements for external CDRMs.

The preceding example, which illustrates a CDRM major node, defines CDRMs that use different dynamic CDRSC options. The resulting capabilities of and restrictions on dynamic cross-domain resource definition are as follows:

- A01M can initiate session requests for any undefined cross-domain resources; however, it can route such requests only to A02M because that is the only other cross-domain resource manager with CDRSC=OPT. A01M can accept session requests from undefined cross-domain resources owned only by A02M because that is the only other cross-domain resource manager with CDRSC=OPT.
- A02M can initiate session requests for any undefined cross-domain resources; however, it can route such requests only to A01M. A02M can accept session requests from undefined cross-domain resources owned by A01M only.
- A03M cannot initiate or accept any session requests for undefined cross-domain resources.

Identifying resources in other domains

This section describes:

- Dynamic definition of cross-domain resources
- Static definition of cross-domain resources
- Model definition of cross-domain resources

Resources controlled by a VTAM in another domain are called cross-domain resources (CDRSCs).

Note: CDRSC definition statements are also used to define independent logical units in the same domain and in nonnative networks. See [“Independent LUs” on page 189](#).

Dynamic definition of cross-domain resources

You do not have to define resources controlled by VTAMs in other domains. VTAM can dynamically create the definition statements to represent resources that reside in other domains. The operands on the CDRM definition statement for the host CDRM and the external CDRM determine whether VTAM can dynamically create CDRSC definition statements.

Notes:

1. You can also dynamically define independent logical units within the same domain using this function, and independent logical units in nonnative networks. See [“Dynamic definition of independent LUs” on page 191](#).

2. To define two or more cross-domain destination LUs that have the same name but are in different networks, use the NQNMODE=NQNAME start option. Otherwise, name conflicts cause session requests involving these LUs to fail.

To have resources in other domains dynamically defined to VTAM:

1. Code your host CDRM with CDRDYN=YES.
2. Code your external CDRM with CDRSC=OPT.

Now all unidentified resources residing under that external CDRM will be dynamically defined to your VTAM (host CDRM). Following is a sample coding for dynamic definition of resources in another domain:

MYHOST	VBUILD CDRM	TYPE=CDRM CDRDYN=YES, SUBAREA=1	X
HOST2	CDRM	CDRSC=OPT, SUBAREA=2	X

Dynamic definition can occur in the following two directions:

- When VTAM receives a session request from an undefined cross-domain resource, the cross-domain resource is the originating logical unit for this session. For example, a terminal in another domain attempts to log on to an application in your domain. The external cross-domain resource manager initiates the session, sending a session request to your VTAM. If the terminal is not defined to your VTAM as a cross-domain resource, VTAM does not accept the session request unless the host CDRM definition is coded with CDRDYN=YES and the external CDRM is coded in your VTAM with CDRSC=OPT.
- When VTAM sends a session request for an undefined cross-domain resource, the cross-domain resource is the destination logical unit for this session. For example, a local terminal attempts to log on to an application in another domain. If the application is not defined as a cross-domain resource, VTAM can still attempt to establish the session by sending a session request to adjacent SSCPs. VTAM sends the session request only if the host CDRM is coded with CDRDYN=YES and the adjacent CDRMs to which VTAM routes the session request are coded with CDRSC=OPT.

When VTAM creates dynamic CDRSCs for destination logical units, it must then dynamically locate the domain in which the resource resides. For this type of dynamic cross-domain resource, VTAM uses a search mechanism to route a cross-domain session request to other adjacent VTAMs. The mechanism is called an adjacent SSCP table and is described in [“Static definition of cross-domain resources”](#) on page 427.

Dynamically defined cross-domain resources are collected in a CDRSC major node named ISTCDRDY. ISTCDRDY is activated automatically during VTAM initialization, and deactivated automatically during VTAM termination.

The dynamic cross-domain resource definition occurs when you have defined the host to allow dynamic definition of CDRSCs.

In general, you have the same control of dynamically defined cross-domain resources as predefined cross-domain resources. You can display and deactivate dynamically defined cross-domain resources using operator commands.

You also generally have the same control of the CDRSC major node ISTCDRDY as other CDRSC major nodes. The DISPLAY ID, MODIFY CDRM, VARY ACT, and VARY INACT commands can all be used for the ISTCDRDY major node.

While the host CDRM is active, the operator can deactivate ISTCDRDY with the VARY INACT command, in which case all dynamically defined CDRSCs are also deactivated and the dynamic CDRSC definition function is disabled. The dynamic CDRSC definition function remains disabled until ISTCDRDY is activated again. You can activate ISTCDRDY again by issuing a VARY ACT command naming ISTCDRDY directly, or by activating the host CDRM (even if the host CDRM is already active). If the host CDRM is inactive, or if the host CDRM is not defined to allow the dynamic CDRSC definition function, activating ISTCDRDY does not provide the dynamic CDRSC definition function.

Dynamically defined CDRSCs are deactivated and deleted by VTAM on a periodic basis if they are not in use, based on the setting of the timer specified in the CDRSCTI start option. A dynamically defined CDRSC

is also deleted if the host that owns that resource fails and the shadow resource for the CDRSC becomes the active definition. For more information about shadow resources, see [“Shadow resources” on page 438](#).

Static definition of cross-domain resources

Cross-domain resources are logical units (application programs, peripheral nodes, and terminals) that are controlled by another VTAM domain. You define cross-domain resources in one or more cross-domain resource major nodes. Each cross-domain resource represents a minor node. You can also define model cross-domain resources in one or more cross-domain resource major nodes. Each model cross-domain resource represents a potential group of similarly named cross-domain resource minor nodes. See [“Model definition of cross-domain resources” on page 428](#) for more information.

You define a cross-domain resource major node with one VBUILD definition statement for the major node, one CDRSC definition statement for each cross-domain resource in the major node, and one CDRSC definition statement for each model cross-domain resource in the major node. The same set of CDRSC definition statements can be used throughout the network.

The name on the CDRSC definition statement represents the name for the resource that is controlled by another VTAM. The CDRM operand, if used, specifies the name of the VTAM that controls that resource. If you do not specify the CDRM operand, the name of the other VTAM that controls the resource is not known. Therefore, you need to use an adjacent SSCP table, so that your VTAM can locate the CDRSC in another VTAM domain. An adjacent SSCP table is a list of other VTAMs with which your VTAM can have an SSCP-SSCP session. This list of VTAMs is used to search your network for the CDRSC. For more information about adjacent SSCP tables, see [“Adjacent SSCPs” on page 431](#).

To statically define two or more CDRSCs that have the same name but are in different networks, use the NQNMODE=NQNAME start option to allow network-qualified names. For more information, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

In the following example of a CDRSC major node, some of the CDRSC definition statements do not specify a CDRM operand, so an adjacent SSCP table is required. However, some of the other cross-domain resource definition statements identify the CDRM owner.

```
*
* EXTERNAL VTAM APPLICATIONS – CDRSCS
*
*          VBUILD  TYPE=CDRSC          CDRSC MAJOR NODE
*
* EXTERNAL VTAM APPLICATIONS – CDRSCS FOR VTAM SUBAREA NODE 03
*
NJE03    CDRSC          JES/NJE
NETV03   CDRSC          NETVIEW
CICS03   CDRSC          CICS
NVAS03   CDRSC          NETVIEW ACCESS SERVICES
*
* EXTERNAL VTAM APPLICATIONS – CDRSCS FOR VTAM SUBAREA NODE 70
*
S04VSCS  CDRSC  CDRM=A04M  VSCS
S04RSCS  CDRSC  CDRM=A04M  RSCS
```

Note: Information about coding the CDRSC definition statement for TSO/VTAM is in [“Multiple-domain network” on page 558](#).

A dependent logical unit and a CDRSC with the same network-qualified name can coexist. In a backup and recovery situation where one host is assuming ownership of a logical unit from another host in the same network, the dependent logical unit can be activated by the new host even though an application program within it currently has a cross-domain session with the logical unit. If the physical and logical units being recovered support ACTPU(ERP) and ACTLU(ERP) requests, the sessions with the logical units are not affected when the physical units and logical units are activated. The CDRSC definition automatically becomes a shadow resource, and the logical unit is now defined as a same-domain (APPL or LU) resource. If the current host wants to relinquish ownership of the logical unit, it releases or deactivates the logical unit and its associated physical unit. VTAM then makes the logical unit a shadow resource, and the CDRSC again becomes the active definition.

Note: For MNPS applications that are active on a network node, you cannot activate a CDRSC definition with the same name. You can activate the CDRSC *before* the MNPS application issuing OPEN ACB, and the CDRSC will become a shadow resource at that time.

When you define your network resources for XRF, consider the following situations:

- If you have predefined all CDRSCs in the network, you need to predefine a CDRSC having the USERVAR name and predefine a CDRSC for both the active and the backup application program.
- If you allow dynamic CDRSCs and trial-and-error rerouting, you do not need to do anything.

Regardless of whether you are using USERVARs with XRF, the names of your USERVARs must be unique. To ensure proper session setup with the desired partner, there can be no other resource within the network that has the same name as any of your USERVARs (except when the USERVAR is identical to its value).

Model definition of cross-domain resources

An alternative method of defining cross-domain resources is with a model CDRSC definition, which can be used as the definition for one or more CDRSCs. You code a model CDRSC definition by placing wildcard characters in the name field of a CDRSC definition statement that defines characteristics for one or more CDRSCs.

If VTAM cannot find an existing definition for a CDRSC, it searches a list of model CDRSC definitions to find the pattern that best matches the name of the cross-domain resource. If VTAM finds a match, it uses that model definition statement to build the definition for the cross-domain resource. If it does not find a match, it attempts to create a dynamic CDRSC, if dynamics are allowed. Model CDRSC definitions enable you to reduce the number of CDRSC definitions you must code in VTAMLST, when dynamic CDRSCs are not an option or are otherwise undesirable.

Model CDRSC definitions can have any keyword specified on them that can be specified on a static CDRSC, which gives you the ability to have different sets of CDRSCs with different characteristics (for example, you might have one model CDRSC coded with DLOGMOD=BATCH and another one coded with DLOGMOD=INTERACT). A CDRSC that is built from a model CDRSC definition is called a clone CDRSC. The clone CDRSC is added to the same CDRSC major node in which the model CDRSC was defined and the clone CDRSC has all the characteristics of the model CDRSC. It is treated as a predefined CDRSC in that any operation that can be done on a predefined CDRSC can be done on a clone CDRSC.

Note: Any information in this publication related to a predefined CDRSC can be assumed to include building a clone CDRSC from an appropriate predefined model CDRSC.

The building of a clone CDRSC is not dependent on the coding of CDRDYN (as a start option or on a CDRM definition statement), nor is it subject to the coding of the CDRSC parameter on the CDRM definition statement. The activation or updating of a model CDRSC does not have any effect on any existing clone CDRSC.

For additional information about modeling cross-domain resources, see [Model definition of VTAM cross-domain resources](#) in *z/OS Communications Server: SNA Resource Definition Samples*.

Coding guidelines

To code a model CDRSC definition, code a CDRSC definition statement to define CDRSC characteristics that you expect to be used by one or more CDRSCs. Use wildcard characters in the name field of the CDRSC definition statement. You can use the following characters:

Asterisk (*)

Represents 0 or more unspecified characters.

Question mark (?)

Represents a single unspecified character.

When placing wildcard characters in the name fields of model CDRSC definitions, you should have some idea of which clone CDRSCs might be built from those model definitions. Use a naming scheme that

ensures that clone CDRSCs are not accidentally built from model CDRSC definitions; that is, ensure that a cross-domain resource does not use a model definition that you did not intend for it to use.

A model CDRSC name, including wildcard characters, can be a maximum of eight characters in length. A question mark (?) can be used anywhere in the model CDRSC name. An asterisk (*) can be used in the second to eighth characters of the model CDRSC name. Model CDRSC names must be unique across all resource names known to this VTAM (including model APPL definitions). There is no defined limit on the number of clone CDRSCs that can be created from one model CDRSC definition, nor is there a defined limit on the number of model CDRSC definitions that you can define.

Model CDRSC definitions can be defined in any number of CDRSC major nodes. Those model CDRSC definitions can appear in a CDRSC major node along with conventionally defined CDRSC definition statements. The model definitions and conventionally defined CDRSC definition statements can appear in any order. If the model is coded before a NETWORK statement with NETID= coded with a valid value, it will be an alias CDRSC and any clones created from it will be considered as predefined alias CDRSCs. If the model is coded after a NETWORK statement with NETID= coded with a valid value, it will be a real CDRSC and any clones created from it will be considered as predefined real CDRSCs.

The rules for the model CDRSC name and how the best matching pattern is found are similar to the rules for model applications. See “How VTAM finds the best match” on page 281 for examples of model names and how matches are found. Table 39 on page 279 gives samples for model names and matching names that apply to CDRSCs and applications.

In addition to the rules for finding the model CDRSC with the best matching pattern for the name, other factors are taken into account when selecting the model CDRSC. These factors are network ID, NQNMODE, and whether a real or alias CDRSC is being requested.

Guidelines: If the input netid for this clone CDRSC request is not the local host netid, but there is another RDTE known by the local host netid, the following apply:

- If an alias CDRSC is being requested, no model is eligible because a clone CDRSC would collide with this RDTE. Therefore, only a dynamic CDRSC is allowed.
- If a real CDRSC is being requested, a model is eligible if the following are true:
 - The model is a real CDRSC.
 - The model's netid matches the input netid.
 - The model has NQNMODE=NQNAME (whether by start option or coded on the model CDRSC).

If the input netid for this clone CDRSC request is the local host netid or if there is no other RDTE known by the local host netid and the input netid for this request is not the local host netid, the following apply:

- If an alias CDRSC is being requested, a model is eligible if one of the following is true:
 - The model is an alias CDRSC.
 - The model's netid is the local host netid.
 - The model has NQNMODE=NAME (whether by start option or coded on the model CDRSC).
- If a real CDRSC is being requested, a model is eligible if the model is an alias CDRSC or if the model's netid matches the input netid.

An example of a CDRSC major node with model CDRSCs follows:

CDRSCSEG	VBUILD	TYPE=CDRSC,...
APPL*	CDRSC	CDRM=HOST1,...
	NETWORK	NETID=NETB
TERM1	CDRSC	CDRM=HOST2,...
APPLB*	CDRSC	CDRM=HOST2,DLOGMOD=BATCH,...
APPLB1*	CDRSC	CDRM=HOST2,DLOGMOD=INTERACT,...
	NETWORK	NETID=NETC
APPLC?	CDRSC	CDRM=HOST3,DLOGMOD=INTERACT,...

In this example, four model CDRSC definitions and one conventional CDRSC definition are shown.

APPL*

A model CDRSC from which alias CDRSCs can be created.

APPLB*

A model CDRSC from which real CDRSCs with a netid of NETB can be created. The clone CDRSCs will have a default logmode of BATCH.

APPLB1*

A model CDRSC from which real CDRSCs with a netid of NETB can be created. The clone CDRSCs will have a default logmode of INTERACT.

APPLC?

A model CDRSC from which real CDRSCs with a netid of NETC can be created. The clone CDRSCs will have a default logmode of INTERACT.

Tips:

- In the above sample major node, if the definition for APPLB* is active, but the definition for APPLB1* is not active, a session request for NETB.APPLB11 creates a clone CDRSC based on the APPLB* definition. That clone CDRSC is used for all session requests for NETB.APPLB11 until the clone CDRSC is deleted, even if the APPLB1* definition, which is a better match, has been activated in the meantime.
- When a clone CDRSC is requested for the origin logical unit of a search, a real CDRSC will be requested.
- When a clone CDRSC is requested for the destination logical unit (DLU) of a search, an alias CDRSC will usually be requested. A real CDRSC will be requested for the DLU if the real name has been learned as a result of alias name translation or a prior host in the search path having a predefined real CDRSC for the DLU.

Note: The name of each model CDRSC must be unique even if different netids are used.

Resource state requirements

A model CDRSC must be active before it can be used to build clone CDRSCs. You can activate a model CDRSC as you activate other CDRSC definitions by doing one of the following actions:

- Issue a VARY ACT command.
- Include the major node in which the model CDRSC is defined in the configuration list that VTAM uses when it is initialized.

By activating the model CDRSC, you ensure that the state of the model CDRSC is active, making it available to build clone CDRSCs.

Clone CDRSCs are different from dynamic (clone) applications in that you can choose to retain the clone CDRSC even when they are deactivated or are no longer being used for sessions. You can do this by specifying DELETE=NO on the model CDRSC definition or on a VARY INACT operator command. The specification of the DELETE operand on the VARY command overrides the specification of the DELETE keyword on the model CDRSC definition. This also means that you can issue VARY ACT against an inactive clone CDRSC, which is not possible for dynamic (clone) applications.

Guidelines:

1. Deleting the clone CDRSCs immediately could result in storage thrashing (freeing and getting clone CDRSCs).
2. Deleting the clone CDRSCs immediately could result in search thrashing (losing search history and then rebuilding it).
3. Not deleting the clone CDRSCs could consume system resources unnecessarily.

For example:

- If a particular set of CDRSCs is located once a day, then you should code DELETE=YES.
- If a particular set of CDRSCs is located many times a day for a short lived session, then you should code DELETE=NO.

Adjacent SSCPs

To supplement the VTAM session request routing service, you can code adjacent SSCP tables or have them dynamically defined. The adjacent SSCP table contains lists of SSCPs that can be in session with a host VTAM and can be used to search for the VTAM that controls a destination CDRSC.

When VTAM receives a session initiation or directory search request for a resource that is not located in that VTAM domain, it attempts to locate the resource by sending the request to its adjacent SSCPs. VTAM routes the request to the SSCP coded in the CDRM operand of the CDRSC definition statement. If one is not coded, the CDRSC is dynamically defined, or the session request with the SSCP in the CDRM operand fails, VTAM uses its adjacent SSCP table.

You can supplement the VTAM search with an adjacent SSCP table even though you code a CDRM operand. VTAM concatenates the proper adjacent SSCP list to the CDRM owner specification and uses the entries for the search. A session request for the CDRSC is routed to each SSCP in the list until the resource is located or the list is exhausted.

An adjacent SSCP table is particularly useful in networks where the SSCP ownership of resources changes frequently. VTAM can locate the owners dynamically rather than by having the programmer maintain the CDRSC definitions. Also, in networks where a large number of cross-domain resources are owned by only a small number of SSCPs, coding an adjacent SSCP table in each host can be easier than coding individual CDRSC definition statements.

If you code adjacent SSCP tables or have them dynamically defined, you do not have to code CDRSC definition statements for logical units in other domains, but VTAM performance is slower because of the time it takes to send startup requests to SSCPs that do not own the logical unit. For information about how to improve the performance of adjacent SSCP tables, see [“Improving performance” on page 435](#).

Alternatively, a coded CDRSC can provide direct session setup, rather than waiting for VTAM to locate a logical unit using the adjacent SSCP tables. By placing CDRSC definitions at a central host (for instance, a CMC host) and placing single entry adjacent SSCP tables in other hosts that point to the central host, session setup requests are sent to the central host if nothing is known about a target resource. When the request reaches the central host, if there is a predefined real CDRSC coded with the resource name, the same NETID as the central host (the CDRSC definitions at the central host must have the same network identifier as the central host), and CDRM equal to the actual owning CDRM (note that the central host does not actually confirm that the coded CDRM is the owner), that information is returned with volatile USERVAR characteristics to the origin SSCP, which can rebuild its adjacent SSCP table with the owning SSCP as the first entry (similar to USERVAR processing). The session setup request can then be forwarded directly to the correct target.

You can also increase control over adjacent SSCP selection by creating adjacent SSCP lists for CDRSCs in an adjacent SSCP table. Then, if an adjacent SSCP list is identified for a CDRSC, session setup requests are sent to only the SSCPs in that list. If the target resource is not owned by (or found through) one of the SSCPs in the list, session establishment fails.

For information about using adjacent SSCP tables for routing requests to other networks, see [“Defining adjacent SSCPs” on page 454](#).

Adjacent SSCP lists for CDRSCs

To define an adjacent SSCP list for a CDRSC, include the ADJLIST operand on the CDRSC definition statement. You can also code the ADJLIST operand on the GROUP definition statement, and the ADJLIST operand sifts down to any CDRSCs in the group for which ADJLIST is not coded. The ADJLIST operand specifies the name of a list of adjacent SSCPs, defined in an adjacent SSCP table using the ADJLIST definition statement followed by any number of ADJCDRM definition statements to identify the specific adjacent SSCPs in the list.

For example, the following sample adjacent SSCP list could be used by A02M to locate LU1 shown in [Figure 128 on page 433](#).

```
A02ADJ      VBUILD TYPE=ADJSSCP
** ADJACENT SSCP LIST: THIS LIST IS USED WHEN THE ADJLIST OPERAND
**                      FOR A CDRSC SPECIFIES LIST1
```

```
**  
LIST1    ADJLIST  
A01M     ADJCDRM
```

See the following sample CDRSC definition statement for LU1.

```
LU1      VBUILD TYPE=CDRSC  
         CDRSC ADJLIST=LIST1
```

When an adjacent SSCP list is identified for a CDRSC, session setup requests are sent to only the SSCPs in the list. The values coded for SORDER and SSCPORD do not affect the search sequence. If the target resource is not owned by (or found through) one of the SSCPs in the list, session establishment fails. This enables strict control over adjacent SSCP selection.

When using adjacent SSCP lists for CDRSCs, the SSCP selection function of the session management exit routine can be used to delete or reorder entries in a list. VTAM cannot dynamically add entries to an adjacent SSCP list.

To add, delete, or change an adjacent SSCP list for a CDRSC, use the MODIFY RESOURCE command. Adjacent SSCP lists for a CDRSC are displayed using the DISPLAY ADJSSCP command. See [z/OS Communications Server: SNA Operation](#) for details on these commands.

Defining adjacent SSCP tables

You can define adjacent SSCP tables by using one, or both, of the following ways:

- Statically define the tables using VTAM definition statements. If this method is used, the SORDER or SSCPORD operand on the NETWORK or CDRM definition statements, or both, can be used to specify a different SORDER or SSCPORD value for each adjacent SSCP table. For more information about SORDER, see [“Using SORDER to control network search order” on page 413](#) or see [z/OS Communications Server: SNA Resource Definition Reference](#). For more information about SSCPORD, see [“Improving performance” on page 435](#) or see [z/OS Communications Server: SNA Resource Definition Reference](#).
- Specify the DYNASSCP start option as YES so that VTAM dynamically builds the tables. The search order (SORDER) for dynamically created adjacent SSCP tables is determined by the SORDER start option.

One or more adjacent SSCP tables can be activated in each VTAM host.

Dynamic definition of adjacent SSCPs

VTAM can dynamically create an adjacent SSCP list. This table consists of all of the adjacent SSCPs with which VTAM has an SSCP-SSCP session. VTAM can then route session-establishment requests to all active adjacent SSCPs until the correct SSCP is found. If you code adjacent SSCP tables and use dynamically defined adjacent SSCP tables, VTAM uses the appropriate adjacent SSCP table without being affected.

To enable dynamically defined adjacent SSCP tables, code the DYNASSCP start option. If you specify YES and VTAM does not locate an appropriate adjacent SSCP table, VTAM dynamically routes the session-establishment request to all active adjacent SSCPs until the correct SSCP is found. If you specify NO, VTAM does not perform this dynamic routing.

By using the dynamic adjacent SSCP table function in conjunction with the adjacent SSCP selection function of the session management exit routine, you can control the order of the dynamic adjacent SSCP table. See [“Improving performance” on page 435](#).

Static definition of adjacent SSCPs

Following is an example of an adjacent SSCP table. This table is coded in host A01M, whose configuration is illustrated in [Figure 128 on page 433](#).

```
A01ADJ    VBUILD TYPE=ADJSSCP  
** DEFAULT SSCP LIST: THIS LIST IS USED WHEN THE CDRM IS  
**           UNKNOWN AND THE NETWORK IS UNKNOWN  
**
```

```

A02M      ADJCDRM      DEFAULT SSCP#1
A03M      ADJCDRM      DEFAULT SSCP#2
**
**
      NETWORK NETID=NETA
**
** ADJACENT SSCPS FOR A CDRM – USED IF CDRSC OWNER IS A02M
**                                AND NETID IS NETA
**
A02M      CDRM
A02M      ADJCDRM      ROUTE ONLY TO THIS HOST
**
** ADJACENT SSCPS FOR A CDRM – USED IF CDRSC OWNER IS A03M
**                                AND NETID IS NETA
**
A03M      CDRM
A03M      ADJCDRM      ROUTE ONLY TO THIS HOST

```

The *cdrmname* used on the adjacent CDRM definition statement in the adjacent SSCP table must be the same as:

- The name on the CDRM definition statement that defines the SSCP to VTAM
- The SSCPNAME used in the VTAM or the name on the host CDRM definition statement in the external VTAM host

Network NETA




Figure 128. Example of adjacent SSCP table connection

In the SSCP of the originating logical unit, the CDRM name of the destination logical unit, if known, must be correct. If the CDRM name of the destination LU is not correct, the operator must issue a MODIFY CDRM command to change the CDRM name.

If SSCPORD=DEFINED and network specific tables have been coded, session requests for DLU resources are not guaranteed to succeed unless at least one entry in the network specific table provides a connection to the DLU.

During adjacent SSCP routing for the DLU resource, SSCP entries are taken from either the network specific table or the default table. If the NETID of the DLU resource is known and a network specific table is coded for that NETID, that table will be used. Otherwise, entries from the default table will be used when the NETID is not known.

It is possible to code an entry in the default table that has a connection to the DLU and not code an equivalent entry in the DLU network specific table. If SSCPORD has the value of DEFINED, a session request for an unknown DLU would be successful because of the working entry in the default table. Subsequent session requests will now use the network specific table because the DLU NETID is known. If the network specific table (matching the NETID of the DLU) does not contain entries which provide a connection to the DLU, the routing for these subsequent sessions will fail.

Table search order

You can individually activate multiple adjacent SSCP table definitions. One or more adjacent SSCP tables can be activated in each VTAM host. Each new set of definitions is added to the internal table used by VTAM. New definitions for destinations already in the internal table replace the current adjacent SSCP list for those destinations, but you cannot deactivate existing adjacent SSCP lists. If a new table replaces an existing one, VTAM issues a message informing the operator.

If the CDRM operand is coded on the CDRSC definition statement, it is used in an adjacent SSCP list. If the CDRM name is an entry in the list, it is moved to the top of the list. If the CDRM name is not an entry in the list, VTAM places the CDRM name at the top of the list. The list is then used to send session setup requests to SSCPs in the list until either the owning SSCP is found or the end of the list is reached. When the owning SSCP is found, VTAM remembers its name. None of these actions affect the actual adjacent SSCP table; they affect only the internal copy of the table that is maintained for each CDRSC.

The adjacent SSCP tables are searched in the following manner:

1. If the NETID and CDRM name are known, a search is done for a table coinciding with the known information.
2. If that table is not found or if only NETID is known, a search is done for a table defined with only the NETID given.
3. If an attempt to find that table fails, or if only the CDRM name is known, the indicated CDRM adjacent SSCP table for all networks is used.
4. If an attempt to find that table fails, the default adjacent SSCP table for all networks is used.
5. If an attempt to find that table fails and if dynamic adjacent SSCP tables are permitted, the dynamic table is used.

Routing with adjacent SSCP tables

If you have defined an SSCP list, VTAM sends the session setup request for an undefined destination logical unit to each SSCP in the list until either the owning SSCP is found or the end of the list is reached. If VTAM does not have an SSCP-SSCP session with an adjacent SSCP in the list, the SSCP is ignored. If a session setup request is routed to an SSCP that does not own the destination logical unit, the request is rejected.

For reliability during session setup, VTAM does the following actions:

- Associates the adjacent SSCP with a cross-domain resource whenever it receives a session initiation request from an SSCP that is not already in the table
- When looking for an SSCP with which to initiate a session, scans the adjacent SSCP table in priority order, giving preference to:
 - SSCPs for which the most recent session initiation attempt succeeded
 - SSCPs for which no session initiation attempts have been made

VTAM bypasses any SSCP in the list with which it does not have an SSCP-SSCP session. If the owning SSCP is not found, the session request is rejected, and the dynamic CDRSC is deleted immediately. VTAM does not wait for the time value specified in the CDRSCTI start option.

Routing rules

Session initiation and INQUIRE APPSTAT requests are routed from the originating LU SSCP to the destination LU SSCP using trial-and-error routing. There are several different means of determining the list of adjacent SSCPs to try when performing trial-and-error routing. However, after the list is determined, the processing is very similar for each of these requests. The rules for routing session requests are as follows:

- The request is not routed if its visit count (SSCP rerouting count) is 0.
- The request is never routed back to the SSCP from which it was received.
- The request is never routed back to its originating SSCP.

- The request is never routed in this network if it was received from an SSCP in this network.
- The request is never routed back to the network from which it was received through the same gateway NCP.

Routing and dynamic tables

If you specify the dynamic adjacent SSCP table function and you have not defined adjacent SSCP tables, VTAM dynamically routes session initiation requests to every SSCP with which it is currently in session. This feature eliminates the requirement to define adjacent SSCP tables. You can, however, continue to define adjacent SSCP tables. VTAM uses the user-defined table if one is defined.

Activating adjacent SSCP tables

Activate adjacent SSCP tables using the VARY ACT command. The ID operand of the VARY ACT command specifies either an adjacent SSCP table or a default SSCP list. Adjacent SSCP tables are not major nodes and cannot be deactivated.

Improving performance

For performance reasons, you might want to change the way VTAM routes session requests. You can use either or both of the following to improve SSCP search performance:

- Use the adjacent SSCP selection function of the session management exit routine to shorten or reorder the list of adjacent SSCPs to which an LU-LU session request is directed. The list is composed of the adjacent SSCPs that are coded in the ADJSSCP table for a specific network or a default list and the list of adjacent SSCPs that was built during the previous LU-LU session setup requests. This exit routine can also be used in conjunction with the dynamic adjacent SSCP facility to avoid having to define any ADJSSCP table while maintaining control of session routing requests.
- Code the SSCPDYN and SSCPORD start options. SSCPORD can be specified as a start option, coded on individual ADJSSCP table definitions, or both. If SSCPORD is not coded for a given ADJSSCP table, the current value of the SSCPORD start option value is used each time that ADJSSCP table is selected for use.

If the CDRM name is coded for the destination CDRSC, that entry is still tried first.

Note: When a resource name that is not valid is entered in a logon request, the request is not rejected until the initiate request has been sent to every SSCP in the default SSCP list and every SSCP has rejected it. The search reduction function can be used to minimize the above.

Adjacent SSCP selection function

VTAM constructs a dynamic adjacent SSCP table in the order in which the CDRM-CDRM sessions are established. Therefore, the search order for a resource might not be the same as that which can be achieved using a statically defined adjacent SSCP table. Extra session establishment time might be required if the SSCP list contains many entries or if it is arranged so that several SSCPs must be tried before the owning SSCP can be located. However, to control the order of search in this case, you can use the session management exit routine (adjacent SSCP selection function) in conjunction with a dynamic adjacent SSCP table.

In the session management exit routine, the adjacent SSCP selection function for LU-LU sessions allows you to modify the list of SSCPs in the search. You can shorten or reorder the list of SSCPs from which the next SSCP used in session setup is chosen during session request routing, or you can decide not to route at all.

The list that the session management exit routine receives as input is either an adjacent SSCP table or a list of adjacent SSCPs that is built during the previous LU-LU session setup for the same resource. If no adjacent SSCPs exist, this function is not invoked. If you do not provide this function in your session management exit routine, VTAM uses the adjacent SSCP table as it exists.

[z/OS Communications Server: SNA Customization](#) describes the adjacent SSCP selection function of the session management exit routine.

SSCPORD and SSCPDYN routing methods

Use the following table to help you decide how to set SSCPDYN and SSCPORD in your system.

Settings	Description
SSCPDYN=YES SSCPORD=PRIORITY	<p>VTAM adds entries dynamically to the adjacent SSCP table. When establishing sessions in releases before VTAM Version 3 Release 2, VTAM gives priority to SSCPs for which the most recent session initiation attempt succeeded or for which no attempt has been made, even if a more direct path to the CDRM that owns the destination LU becomes available.</p> <p>When establishing sessions, VTAM Version 3 Release 2 and later releases give priority to the CDRM that owns the destination LU (if known), and then to SSCPs for which the most recent session attempt succeeded.</p> <p>This combination of options gives you the greatest flexibility for setting up routes across networks, and, if your adjacent SSCP table is large, it gives you the best performance during session setup.</p>
SSCPDYN=YES SSCPORD=DEFINED	<p>Entries are added dynamically to the adjacent SSCP table, but VTAM always searches the table from top to bottom when it tries to initiate a session.</p> <p>Because the dynamic entries are added at the bottom of the list, they are tried last during session initiation.</p> <p>This combination of options results in increased session setup time if the primary paths (those near the top of the list) are not available.</p> <p>But if a secondary path (one nearer the bottom of the table) is used, it does not get priority for subsequent sessions.</p>
SSCPDYN=NO SSCPORD=PRIORITY	<p>No entries are added to the adjacent SSCP table, except when the owner of a cross-domain resource is defined explicitly by the CDRM operand of the CDRSC definition statement.</p> <p>Therefore, VTAM can use only the paths you define in the table, and if one of these paths fails, VTAM does not attempt to use it again until all other paths also fail.</p> <p>This combination of options tends to limit the number of available cross-network routes.</p> <p>You might also need to define additional adjacent SSCP tables because VTAM does not build them automatically each time an application program issues a CLSDST PASS macroinstruction for a cross-network session.</p>

Settings	Description
SSCPDYN=NO SSCPORD=DEFINED	<p>No entries are added to the adjacent SSCP table, except when the owner of a cross-domain resource is defined explicitly by the CDRM operand of the CDRSC definition statement. VTAM always searches the table from top to bottom when it tries to initiate a session. Again, VTAM can only use the paths you define in the table; but, because VTAM always searches the table from top to bottom, a path that fails can be used again as soon as it is restored.</p> <p>This combination of options results in increased session setup time if the primary paths (those near the top of the list) are not available.</p> <p>But if a secondary path (one nearer the bottom of the table) is used, it does not get priority for subsequent sessions.</p> <p>You might also need to define additional adjacent SSCP tables because VTAM does not build them automatically each time an application program issues a CLSDST PASS macroinstruction for a cross-network session. VTAM does not add new entries unless it is the owner of the resource. The owner is determined by the CDRM operand of the CDRSC definition statement, or by session initiation.</p>

CDRM owner verification for cross-domain resources

You can enable SSCPs to take over cross-domain resources without requiring intermediate SSCPs to issue a MODIFY CDRM command. This enables VTAM to automatically change the CDRM ownership of a CDRSC. This can be useful in cross-network sessions where one of the session partners has been taken over by a different SSCP and an attempt is made to establish a new session.

Note: Owner verification of a cross-domain resource is optional.

This SSCP takeover of cross-domain resources is accomplished in the definition of the CDRSC with the VFYOWNER operand. VFYOWNER=YES can be specified only for a CDRSC that is coded with the NETID and CDRM operands. Specifying VFYOWNER=NO (which is the default) allows VTAM to automatically change the CDRM ownership of a CDRSC. VFYOWNER=YES in the CDRSC definition makes VTAM reject session setup requests that contain a conflicting owner. The default (VFYOWNER=NO) allows the takeover of a CDRSC without the need for intermediate SSCPs to issue a MODIFY CDRM command.

The MODIFY CDRM command does not need to be used for dynamically defined CDRSCs. If a session setup request is received for a dynamically defined CDRSC from a CDRM other than the current CDRM owner, VTAM automatically changes the CDRM ownership of the CDRSC. However, the operator might still want to use the MODIFY CDRM command if LU-LU session recovery is initiated by the logical unit that received the original session request.

For example, a terminal user logs on to an application program, and the CDRSC representing the terminal is defined dynamically in the domain of the application program. If the session between the application program and the terminal is disrupted, the application program might try to reinitiate the session instead of requiring the terminal user to log on again. In this case, the VTAM operator in the domain of the application program needs to use the MODIFY CDRM command to change the CDRM owner of the CDRSC representing the terminal, if the terminal is taken over by a backup SSCP other than the SSCP in the application program domain.

A VTAM domain that takes over a cross-domain resource does not need to use the MODIFY CDRM command because in this case, the VTAM shadow resource function applies. For more information about shadow resources, see [“Shadow resources” on page 438](#).

Changing ownership of cross-domain resources

Each cross-domain resource (CDRSC) definition statement indicates its controlling SSCP (CDRM=*cdrmname*). The controlling SSCP is also called the owning CDRM of the cross-domain resource. The ownership of a cross-domain resource can be changed if the owning CDRM loses its SSCP-LU session with the logical unit, and another SSCP takes over the logical unit.

You can use the MODIFY CDRM command to inform VTAM of the takeover, naming the new owner of the cross-domain resource. This command can be used whether the cross-domain resource is active or inactive. If the cross-domain resource is already in session with a logical unit, the session is undisturbed and the owner is saved until the session terminates. Session startup requests continue to be routed to the old owner. Specifying TYPE=IMMED on the MODIFY CDRM command can be used to force the ownership change immediately regardless of active sessions.

Shadow resources

The VTAM shadow resource function enables VTAM to dynamically change what it defined previously as CDRSCs into same-domain resources during takeover.

A shadow resource is an alternate definition of a network resource. For example, an LU is defined to a backup host both as a same-domain resource and as a cross-domain resource. Under normal operation, the LU is a cross-domain resource, and the cross-domain definition is used. The same-domain resource definition, which is not used, is the shadow resource. Whichever definition is not being used is considered to be the shadow resource.

If the host that owns the LU fails, the backup host assumes ownership of the LU, the same-domain definition becomes the active definition, and the cross-domain definition becomes the shadow resource. Resource takeover can be accomplished without disrupting existing active sessions. When the original host becomes available, the backup host relinquishes ownership of the LU. The backup host then activates its cross-domain definition for the LU, and its same-domain definition once again is the shadow resource.

When the LU is owned by the backup host, other hosts that need to access the LU must be able to communicate with and route requests to the backup host.

Chapter 18. Connecting multiple subarea networks

This section describes a multiple-network environment that uses SNA network interconnection (SNI). A multiple-network environment consists of multiple independent SNA networks that are interconnected. The SNI facility enables communication between these separate networks.

SNA network interconnection enables you to:

- Interconnect multiple, independent SNA networks so that terminal users in one SNA network can access information (application programs) in other SNA networks
- Divide an existing network into smaller, independent networks, with communication between these networks maintained by the SNI function

Following are descriptions of environments in which SNI can be used. (An enterprise is a business organization.)

Independent enterprises

If two or more independent enterprises need to access and exchange information using their current SNA telecommunication networks, SNI can be used to interconnect these autonomous networks, which might have dissimilar characteristics.

Enterprise mergers and acquisitions

If two enterprises have merged or if one enterprise has acquired another, and the dissimilar SNA networks need to interconnect, SNI can be used so that any terminal or application program in one network can access any terminal or application program in the other network.

Multiple networks within an enterprise

If multiple SNA networks have evolved within a large enterprise, each with different network characteristics, SNI can be used to merge these dissimilar networks into a single logical network, yet maintain the autonomy of each.

In each of the preceding examples, SNI can be used to interconnect networks with dissimilar characteristics while preserving the independence of each participating network. These participating networks do not have to change to a common address structure or resource-naming convention, and each network can preserve its existing management procedures and controls.

Defining a multiple-network environment

Defining your multiple-network environment involves identifying your different networks and defining the resources that will help you communicate between the networks. You need the following for controlling cross-network sessions:

Gateway VTAMs

Enable at least one gateway VTAM for each network gateway. The gateway VTAM controls sessions across the network gateway. The gateway VTAM is in session with the gateway NCP, assisting in address translation and cross-network routing. The gateway VTAM must be either an MVS host running VTAM Version 2 Release 2 or higher releases or a VM host running VTAM Version 3 Release 1.1 or higher releases.

Gateway NCPs

You must specify at least one gateway NCP for each network gateway. The gateway NCP works with the gateway VTAM to control address translation for cross-network sessions. The gateway NCP provides network address translation from a pool of network alias addresses for each network. The gateway NCP also assists in activating explicit and virtual routes for cross-network sessions.

The gateway NCP is a resource in each network to which it connects. Each gateway NCP has one native network, where the gateway VTAM that can activate the NCP subordinate resources resides. Each gateway NCP also has one or more nonnative networks. For example, in [Figure 129 on page 440](#), NCP21 is subarea 21 in NETA and subarea 51 in NETX. NCP21 native network is NETA; its nonnative network is NETX.

The configuration in [Figure 126 on page 421](#) is the same as the configuration in [Figure 129 on page 440](#), except that [Figure 129 on page 440](#) shows a multiple-network environment because it contains two independent SNA networks. NCP21 in NETA is also subarea 51 in NETX and NCP11 in NETB is also subarea 52 in NETX. There are two NCP11s in this environment. The names can be the same because they reside in different networks (hosts, however, cannot have the same name because of CDRM conflicts).

Code everything that you would for [Figure 126 on page 421](#). NETID must be coded in the start options to the definition of HOST1 in NETA:

Add at least the following to the definition of HOST2 in NETA:

- GWSSCP=YES must be defaulted or coded and NETID should be coded in the start options.
- One NETWORK definition statement for each CDRM or group of CDRMs, identifying which CDRMs reside in each network.
- CDRM definition statement in the cross-domain resource manager major node for VTAM3 in NETB.
- A GWPATH definition statement following the cross-domain resource manager minor node for VTAM3 in NETB. Code the ADJNET, ADJNETEL, ADJSA, and GWN or SA operands.
- NETID and GWCTL operands on the PCCU definition statement defining NCP21 in NETA.
- NETID, HSBPOOL, and CANETID operands on the BUILD definition statement in the NCP major node for NCP21 in NETA.
- GWNAU definition statement in the NCP major node to identify subarea 52 in NETX.
- NETWORK definition statement in the NCP major node to identify NETX. Code the NUMHSAS and SUBAREA operands.
- GWNAU definition statement to represent NCP21 in NETA as NCP51 in NETX.

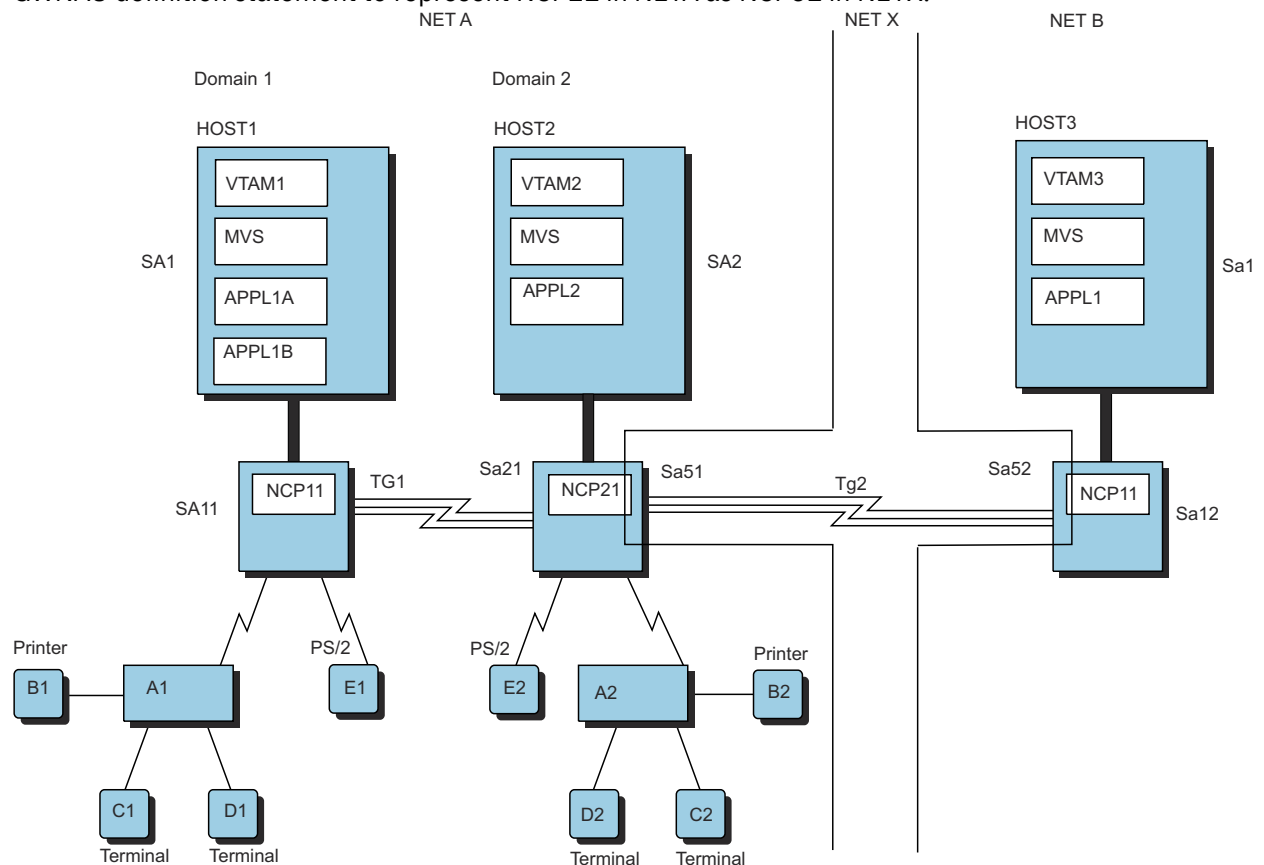


Figure 129. Multiple-network environment

Also code at least the following in HOST3 in NETB:

- Start options for HOST3 in NETB (NETID, SSCPID, and SSCPNAME); code or default to GWSSCP=YES.

- Application program minor node for APPL1 in NETB.
- One NCP major node for NCP11 in NETB.
- One NETWORK definition statement for each CDRM or group of CDRMs, identifying which CDRMs reside in each network.
- CDRM definition statements in the cross-domain resource manager major node for VTAM3 in NETB and VTAM2 in NETA.
- A GWPATH definition statement following the cross-domain resource manager minor node for VTAM2 in NETA. Code the ADJNET, ADJNETEL, ADJSA, and GWN or SA operands.
- NETID and GWCTL operands on the PCCU definition statement defining NCP11 in NETB.
- NETID, HSBPOOL, and CANETID operands on the BUILD statement in the NCP major node for NCP11 in NETB.
- GWNAU definition statement in the NCP major node to identify NCP51 in NETX.
- NETWORK definition statement in the NCP major node to identify NETX. Code the NUMHSAS and SUBAREA operands.
- GWNAU definition statement to represent NCP11 in NETB as NCP52 in NETX.

Note: Resources with identical names can exist in interconnected networks. To resolve naming conflicts, use the NetView alias name translation facility or the NQNMODE=NQNAME start option. For information about the NetView alias name translation facility, see [“NetView alias name translation facility” on page 466](#). For information about the NQNMODE=NQNAME start option, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

SNA network interconnection provides the following major functions:

Cross-network session control

Enables VTAM to perform cross-network session initiation and termination. As required, VTAM can request address translation, name translation, and configuration data.

Address translation

Enables networks with different addressing schemes to communicate. The gateway NCP performs the function of address translation when crossing network boundaries.

Resource name translation

Enables networks with conflicting resource names to communicate. This function can be performed by the NetView program or the VTAM session management exit routine.

Note: The NQNMODE=NQNAME start option is an alternative to resource name translation. For more information, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Cross-network problem determination and network management

Collects and correlates problem determination and network configuration information for cross-network sessions. The NetView program provides the function of cross-network problem determination and network management.

SNI configurations

This section describes some basic SNI configurations and their attributes.

Single gateway configuration

This configuration consists of one gateway VTAM and one gateway NCP. This configuration is the easiest to configure and define, and cross-network session control is centralized. However, failure of the gateway NCP disrupts all cross-network sessions. This type of gateway can be used to connect up to 255 networks, with the gateway NCP residing in multiple nonnative networks.

Single gateway VTAM with multiple gateway NCPs

In this configuration, two or more NCPs act as the gateway NCPs. This configuration has the advantages of the single gateway configuration, but the failure of one gateway NCP does not disrupt all cross-network sessions.

Serial configuration

In this configuration, one gateway VTAM resides in a transport network and controls two or more gateway NCPs that connect to the nonnative networks. This configuration is best if the networks on either side of the transport network do not exchange much information.

Back-to-back configuration

This configuration consists of two gateway VTAMs and two gateway NCPs that connect through a null network. This configuration provides the most network isolation and security, but it requires more network definition. It also means that address and name translation must be performed twice, and the cross-network session path is longer. The SSCP-SSCP session between HOST2 and HOST3 in [Figure 129 on page 440](#) is an example of a back-to-back configuration.

Start options defining other networks

The start options that you set for VTAM can determine whether the VTAM is a gateway VTAM or a nongateway VTAM.

Note: You can use the XNETALS start option to connect to nonnative type 2.1 PUs. For more information see [“Nonnative network type 2.1 connections” on page 184](#).

Start options for gateway VTAMs

To create a gateway VTAM, specify the following start options:

NETID (required)

NETID specifies the name of the network that contains this gateway VTAM. The network identifier allows VTAM to determine which gateway NCP definition statements apply to the gateway host.

SSCPID (required)

Change SSCPID as necessary. SSCPID uniqueness must be maintained across connecting networks because two SSCPs with the same ID cannot establish a session.

SSCPNAME (required)

SSCPNAME specifies the unique name by which this gateway SSCP is known to the gateway NCP and other networks. This name must be the same as the NAME operand on a GWNAU definition statement that is used to predefine an SSCP element address in a gateway NCP. This name must also be the same as the owner name (CDRM operand) used on the CDRSC definition statements in other networks that represent resources owned by this VTAM.

HOSTPU (recommended if the NetView program is installed)

HOSTPU specifies a unique network name for the VTAM host physical unit. The NetView program uses this name to determine which VTAM host physical unit it is tracing. (If HOSTPU is not specified, all VTAM host physical units are named ISTPUS.)

GWSSCP=YES (required)

GWSSCP=YES indicates that this SSCP can serve as a gateway SSCP.

Note: Even if this SSCP is not performing a gateway function, GWSSCP can be used to enable this SSCP to act as an intermediate SSCP during a session setup.

GWSSCP start option for nongateway VTAMs

If you code GWSSCP=NO, trial-and-error routing of session initiation requests in the originating domain can be done, but no rerouting is possible for session initiation requests received by this host from other SSCPs.

Coding GWSSCP=YES for a VTAM that does not activate gateway NCPs permits VTAM to reroute received session establishment and directory requests. It also enables you to code GWPATH definition statements in conjunction with a CDRM definition to define multiple gateway paths to that CDRM.

Configuration lists for gateway VTAMs

You might want to update the initial configuration list to include any new major nodes and tables, such as:

- Gateway NCP major nodes
- CDRM major nodes
- CDRSC major nodes
- Adjacent SSCP tables
- Path tables

Connecting networks

When you are planning to connect networks, you should consider the following, and if necessary, discuss these things with the system programmers for the networks that are being connected:

- To guarantee uniqueness, know the SSCPIDs and the NetView domain IDs used in each network, or use a naming convention.
- If you own a gateway NCP:
 - Obtain the MAXSUBA specification used for any attaching network that contains nonextended subarea addressing nodes so that you can code the MAXSUBA operand on the NETWORK definition statement in the gateway NCP.
 - Specify an available subarea number in the attaching network on the SUBAREA operand on the NETWORK definition statement.
 - Know how many subarea nodes can communicate with the gateway NCP in the attaching network so that you can code the NUMHSAS operand on the NETWORK definition statement in the gateway NCP.
 - Exchange information about explicit route and virtual route structures so that you can code PATH definition statements for nonnative networks in the gateway NCP.
- If you plan to code GWPATH definition statements, you need to know the subarea and element addresses of all CDRMs in adjacent networks with which your gateway VTAM is to establish sessions.
- If you are not using network-qualified names, to avoid future name conflicts, decide upon a naming convention for alias names and for real names.
- If you want to guarantee that certain application programs can always obtain an alias address, you need to know the real logical unit names of those application programs. These names are required on the GWNAU definition statements that predefine specific resources in the gateway NCP.
- Tell system programmers for the other networks the names to be used for logical units. Discuss equivalent COS and logon mode entry names.
- Some application programs require a definition for each of their session partners. If these application programs are to have cross-network sessions, define one of the following to the application programs:
 - Alias names for session partners in other networks
 - LU alias names (on the CDRSC definition statements) for session partners in other networks
 - USERVARs that map to the network-qualified names of session partners in other networks
- If you code a CDRM operand on a CDRSC definition statement and you specify a NETWORK definition statement for it, use the real SSCP name for the host in the other network.

The real SSCP name is the name specified in the SSCPNAME start option. For any host for which SSCPNAME is not specified, the real SSCP name is the name of the host CDRM for that host. You should use this name on your definition for this CDRM if you code such a definition.

- If you code a CDRM definition statement in an adjacent SSCP table, specify the real SSCP name on the label of the CDRM definition statement.
- If you code an SSCP name on an LU definition statement for the NetView alias name translation facility, you must specify the real SSCP name.
- You must code the real SSCP name on the ADJCDRM operand in an adjacent SSCP table.
- Ensure that the COS table named in the COSTAB operand of the BUILD definition statement or a NETWORK definition statement in the gateway NCP is available to the gateway SSCP. You must install it in VTAMLIB. The table specified in the COSTAB operand is used for mapping Class of Service table

entries for the routes that originate in the gateway NCP and terminate in the network being defined. If a COS table for the connected network is not supplied, the COS table for your network must include any COS names that can be requested by logical units in other networks.

- Ensure that NCP data transfer related operands (such as UNITSZ, TRANSFR, and BFRS) that are coordinated between domains in a multiple-domain network are also coordinated between interconnected networks.
- If you use the NetView program, you need to know the NetView suppression characters used in every domain in every network to ensure that the characters are all the same.

Defining a gateway VTAM

The gateway VTAM, in conjunction with the gateway NCP, performs cross-network session initiation and termination by acting as an intermediary between VTAMs in the connected networks. (After the session is set up, there is no need for the LU-LU session path to go through the gateway VTAM.) The gateway VTAM also directs the gateway NCP to set up the address transformations and routes needed for cross-network sessions. The gateway VTAM determines the routing for cross-network sessions and resolves resource names in conjunction with the name translation facility.

Note: A gateway VTAM must be either an MVS host with VTAM Version 2 Release 2 or higher, or a VM host with VTAM Version 3 Release 1.1 or higher.

The gateway VTAM can be in any one of the connected networks.

To convert a non-SNI VTAM to a gateway VTAM, you need to make the following changes:

- Update your start options.
- Update your initial configuration list.
- Update your application program major node if you are using NetView alias name translation facility.
- Update your CDRM major node.
- Create CDRSC definition statements (optional).
- Create an adjacent SSCP table.

In a gateway configuration, VTAM has several different names. The following three names are assigned using start options:

- The ID specified on the SSCPID start option
- The name specified on the SSCPNAME start option
- The name specified on the HOSTPU start option

These are the names by which this SSCP CDRM is known by other domains in this and other networks. No requirement exists for these names to match any of the other names listed previously. However, you might want to have each CDRM known by a single name throughout all the networks. You can do this by coding a single set of CDRM definitions that is used in all the connected networks. Coding the same name for this SSCP on all definitions makes network management easier.

Another name can be assigned on the CDRM definition statement. However, if that name differs from the SSCPNAME that is used, the results are unpredictable because of VTAM unique use of each of the varying names. It is recommended that these names be the same to avoid inconsistencies. One exception to this is that HOSTPU must be different from the CDRM name or SSCPNAME.

Defining cross-domain resource managers

You need to code CDRM definition statements for VTAMs in other networks with which this VTAM requires direct communication.

Procedure

Perform the following steps in a CDRM major node:

1. Add one NETWORK definition statement for each network to be connected through a gateway NCP. This is necessary to identify the networks that contain the defined CDRMs.

The NETID operand of the NETWORK definition statement identifies the network that contains the defined CDRMs. All the CDRMs that follow a given NETWORK definition statement must belong to that network.

The value coded for NETID on the NETWORK definition statement must match the NETID start option specified for the CDRM in its own domain. CDRM definition statements for CDRMs in the same network as this host CDRM can follow a NETWORK definition statement whose NETID identifies the network that contains this host CDRM.

By coding a NETWORK definition statement for every network (including that of the host CDRM), you can use the same set of CDRM definitions at hosts in different networks. (However, these hosts must all have VTAM Version 2 Release 2 or later releases.)

2. Add one CDRM definition statement for each SSCP in another network with which the gateway VTAM is in session (this is required).

The CDRM names defined in the gateway SSCP must be unique and must immediately follow their own NETWORK definition statement.

3. Add one or more GWPATH definition statements after each CDRM definition statement that defines an SSCP in another network.

If you specify the GWSSCP=YES start option, you can use a GWPATH definition statement in conjunction with a CDRM definition statement. In this case, the value of the ELEMENT operand on the GWPATH definition statement must match that defined with a GWNAU definition statement in the gateway NCP through which the gateway SSCP is reached. If the gateway SSCP can be reached through multiple gateway NCPs, multiple GWPATH definition statements can be associated with the single CDRM definition statement defining the gateway SSCP. (This is required for CDRMs in other networks that are not predefined in the gateway NCP.)

GWPATH definition statements define one or more cross-network session paths to be used for cross-network LU-LU and SSCP-SSCP sessions.

If no GWPATH definition statement is coded, VTAM uses the SUBAREA value specified on the CDRM definition statement to select a gateway NCP. The ELEMENT value on the CDRM definition statement must match that specified by the GWNAU definition statement for a predefined CDRM in the gateway NCP.

The ADJNETSA and ADJNETEL operands on the GWPATH definition statement must specify the subarea and element numbers of the CDRM as they are known in the adjacent network. For a back-to-back gateway NCP configuration, these numbers are obtained from the SUBAREA and ELEMENT operands on the GWNAU definition statement that defines this external CDRM in the second gateway NCP in the back-to-back configuration. For information about how the gateway path is selected, see [“Gateway path selection” on page 460](#).

Using the ADJNETCS operand on the GWPATH definition statement, you can specify the Class of Service name to be used for establishing the route for the SSCP-SSCP session. Usually, this Class of Service name is ISTVTCOS. If you specify ADJNETCS, VTAM does not invoke the NetView alias name translation facility or the session management exit routine to perform Class of Service translation.

4. The value you code on the SUBAREA operand of the CDRM definition statement should be the subarea in the gateway NCP through which the gateway SSCP is reached.

The ELEMENT operand for the CDRM definition statement must match the element predefined for the gateway VTAM with a GWNAU definition statement in that gateway NCP. If these operands are not defined, it is impossible to initiate an SSCP-SSCP session between the nongateway VTAM and a gateway SSCP in another network.

If the gateway SSCP can be reached through multiple gateway NCPs, it appears as multiple CDRMs (that is, as SSCPs in different subareas) from the point of view of the nongateway SSCP. If a gateway NCP failure disrupts the SSCP-SSCP session between the nongateway SSCP and the gateway SSCP, reactivation of the session through a different gateway NCP is possible if you provide multiple CDRM definition statements for the gateway SSCP. Each of these definition statements must have a unique

label. The SUBAREA operand specifies the subarea of one of the gateway NCPs, and the element value is reserved for the gateway SSCP with a GWNAU definition statement coded in that gateway NCP subarea. Given these definitions, the operator at either the gateway SSCP or the nongateway SSCP can reestablish the SSCP-SSCP session through an alternate gateway NCP.

Results

Figure 130 on page 446 is an example of an SNI back-to-back configuration.

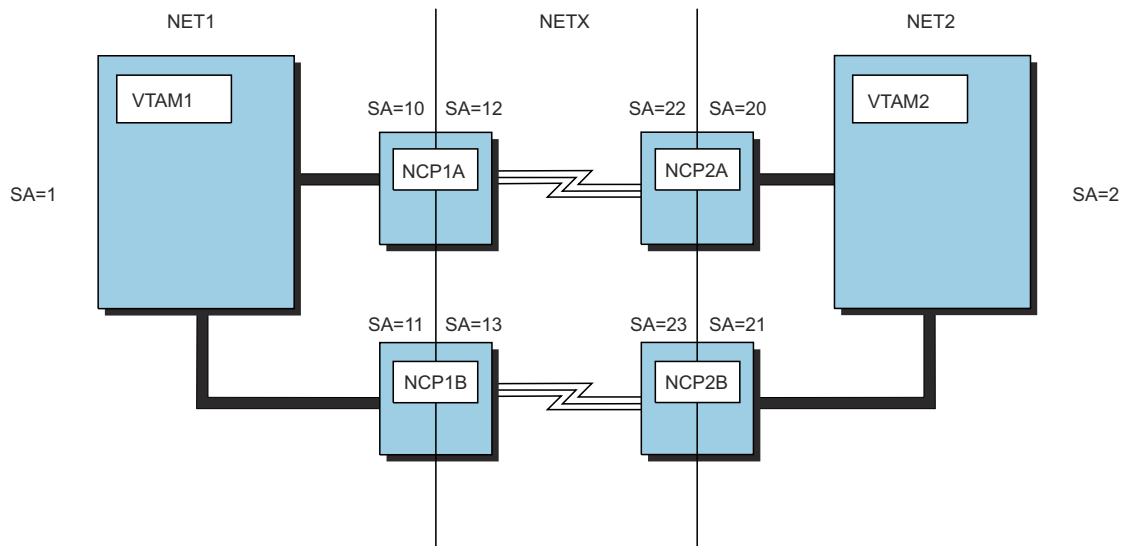


Figure 130. Multiple-network configuration: CDRM major nodes

Code the following in the CDRM major node in VTAM1:

```
NET1    NETWORK NETID=NET1
VTAM1   CDRM  SUBAREA=1,
:
NET2    NETWORK NETID=NET2
VTAM2   CDRM
NCP1A   GWPATH SUBAREA=10,
        ELEMENT=1,
        ADJNET=NETX,
        ADJNETSA=22,
        ADJNETEL=1,
NCP1B   GWPATH SUBAREA=11,
        ELEMENT=1,
        ADJNET=NETX,
        ADJNETSA=23,
        ADJNETEL=1,
```

It is recommended that you code the ELEMENT operand rather than accepting the default. In this way, it is easier to ensure that you do not have the same subarea and element number for multiple resources.

The CDRM major node in VTAM1 shows that you need to code a GWPATH definition statement for each gateway NCP. Note that you need to be careful to match the NCP pairs (NCP1A with NCP2A and NCP1B with NCP2B) to ensure that the CDRM-CDRM session uses the same pair of gateway NCPs.

Thus, code the following in the CDRM major node in VTAM2:

```
NET2    NETWORK NETID=NET2
VTAM2   CDRM  SUBAREA=2,
:
NET1    NETWORK NETID=NET1
VTAM1   CDRM
NCP2A   GWPATH SUBAREA=20,
        ELEMENT=1,
        ADJNET=NETX,
        ADJNETSA=12,
        ADJNETEL=1,
NCP2B   GWPATH SUBAREA=21,
        ELEMENT=1,
```

```
ADJNET=NETX,  
ADJNETSA=13,  
ADJNETEL=1,
```

Defining cross-domain resources

This topic introduces methods of defining resources in other networks.

- Use a CDRSC definition statement (or a model CDRSC definition statement) without an owning network specified.

This method of defining CDRSCs can be used by hosts within either of the networks in which the session partners reside or by intermediate networks. The definition can be used to satisfy either a real or an alias name, but if the NetView alias name translation facility or the alias function in the session management exit is used, you cannot define a CDRSC for both. For example, in HOST1 (in [Figure 132 on page 450](#)), you can code the following statements:

```
CDRSCSEG  VBUILD      TYPE=CDRSC, ...  
TERM1     CDRSC       CDRM=HOST2, ...
```

To code the NQNMODE or LUALIAS operand for a CDRSC, the owning network must be specified.

- Use a CDRSC definition statement (or a model CDRSC definition statement) with an owning network specified.

This method of defining CDRSCs can be used by any of the hosts involved in establishing the requested session. This method provides for the best session setup performance. For example, in HOST1 (in [Figure 132 on page 450](#)), you can code the following statements:

```
CDRSCSEG  VBUILD      TYPE=CDRSC, ...  
TERM1     NETWORK    NETID=NETB  
          CDRSC       CDRM=HOST2, ...
```

This is the method to use if you want to define cross-domain resources by their network-qualified names. To define cross-network resources by only their network-qualified names, specify NQNMODE=NQNAME (in the definition or using the start option). To define cross-network resources by both their network-qualified names and their unqualified names, specify NQNMODE=NAME (in the definition or using the start option).

The LUALIAS operand is valid only for cross-network resources with an owning network specified.

- Use a dynamically defined CDRSC.

Instead of predefining CDRSCs, you can decide whether some or all CDRSCs are dynamically allocated and defined. The host that owns the logical unit initiating a session can dynamically define the requested session partner. In the intermediate SSCP, both the originating LU and the destination LU can be dynamically defined. In the destination LU SSCP, the originating LU can be dynamically defined. Independent LUs acting as either the destination or the originating LU can be dynamically defined at either the destination or originating LU SSCP.

Dynamically defined CDRSCs take the NQNMODE value that is defined on the NQNMODE start option. The NQNMODE start option affects dynamically defined CDRSCs in the following way:


- If NQNMODE=NAME, a cross-network dynamic CDRSC is defined by its network-qualified name and by its alias name and alias network, even if the alias name and alias network are assumed (that is, the alias name is equal to the real name and the alias network is equal to the network of the session partner).
- If NQNMODE=NQNAME, a cross-network dynamic CDRSC is defined by its network-qualified name. It can also have a dynamic LUALIAS created for it if the session partner application specified LUAPFX on its APPL definition statement. A cross-network dynamic CDRSC is not defined by its assumed alias name and alias network unless the name portion of the alias name and the real name are different.

Session initiation request

In Figure 131 on page 448, assume that HOST11 and HOST21 are gateway VTAMs, and APPL1 wants to initiate a session with TERM1. The value of NQNMODE, either as a start option or defined on the CDRSC definition statement, affects the name by which the cross-network resource is known. If NQNMODE=NAME is specified as a start option or on the CDRSC definition statement for TERM1, because there is no TERM1 in NET1, TERM1 real name is TERM1 in NET3 and its alias name is TERM1 in NET1. (If a resource TERM1 actually existed in NET1, TERM1 in NET3 must be translated to a name that is not being used in NET1, such as TERM3 in NET1.)

If NQNMODE=NQNAME is specified as a start option or on the CDRSC definition statement for TERM1, TERM1 real name is NET3.TERM1. It is *not* known by the alias name NET1.TERM1. (If a resource TERM1 actually existed in NET1, and the session initiation request used a nonnetwork-qualified name, the session would be established with TERM1 in NET1. If the session initiation request used NET3.TERM1, the session would be established with TERM1 in NET3. Another option to using a network-qualified name in a session request would be to translate TERM1 to a name that is not being used in NET1, such as TERM3 in NET1.)

Note: A network-qualified session initiation request (for example, NET3.TERM1) cannot use the application ACBNAME unless it is the same as its APPL name.



Note: The connections between host and NCPs have been omitted for clarity. You can assume that each NCP connects to each host in the networks to the left and right of it.

Figure 131. Example of three interconnected networks

The session initiation request that is flowing from HOST11 to HOST21 contains the following information:

- The alias name only (TERM1 in NET1), if APPL1 initiates a session for TERM1 with a nonnetwork-qualified name, and the CDRSC for TERM1 is dynamically defined or predefined using an alias name in HOST11.
- Both the alias name (TERM1 in NET1) and the real name (TERM1 in NET3), if APPL1 initiates a session for TERM1 using a network-qualified name (NET3.TERM1), or if the session initiation request is network qualified.

The request flow described here assumes that GWCTL=SHR is in the host PCCU definition statement. If, according to the rules for gateway control, HOST11 is designated to request an alias address for GWN1, the real name for TERM1 must be properly determined in HOST11 by defining a CDRSC with NETWORK specified or by using the NetView alias name translation facility.

Name assumption

In [Figure 131 on page 448](#), assume that HOST21 is a gateway VTAM. Also assume that HOST11 sends only the alias name to HOST21 in the session initiation request. The following steps occur:

- HOST21 dynamically defines a CDRSC for TERM1 or uses a predefined CDRSC.
- HOST21 attempts to determine the real name of TERM1 using the NetView alias name translation facility or the session management exit routine.
- HOST21 routes the session setup request to HOST31 using the adjacent SSCP tables, and the request leaves the first gateway.
- When the request leaves the first gateway, VTAM does the following actions:
 - If the real name and network identifier are not determined using the NetView alias name translation facility or the VTAM session management exit routine, VTAM sends out a search using the alias name to find the real name and network identifier.
 - The alias CDRSC representation (TERM1 in NET1) is updated to reflect the real CDRSC representation (TERM1 in NET3).
- In SSCP21, a CDRSC is defined or dynamically created for APPL1 also. Because APPL1 is the originating LU, its real name and network identifier are known. The predefined CDRSC can be defined with or without a network identifier.
- HOST31 receives the session initiation request from HOST21, which contains both real and alias names for TERM1. HOST31 then determines that it owns TERM1.
- HOST31 locates a CDRSC representation for APPL1 or dynamically defines one using the real name (APPL1 in NET1). If you have defined a CDRSC definition statement for APPL1 in NET3, HOST31 will use that CDRSC definition statement and will not use the dynamically defined CDRSC.
- When the CDRSC is found, HOST31 merges any pertinent information from the alias CDRSC representation with the real representation.
- If session initiation request processing is successful, HOST31 sends a positive response.
- HOST11 receives the response, and further processing depends on how TERM1 is defined in HOST11.
 - If TERM1 is dynamically allocated, VTAM merges the acquired information (real network ID and owning SSCP name) into the existing CDRSC.
 - If TERM1 is predefined using its alias name, the predefined CDRSC representation is updated with the real name, network identifier, and owning SSCP.
 - If TERM1 is predefined using its real name and network identifier, no further processing is required. Session setup continues, and resources are recognized by their real names, alias names, and network identifiers.

If both the real and alias names are passed in the session initiation request, either a predefined real CDRSC representation for TERM1 in NET3 is located, or a dynamic CDRSC is allocated. A predefined CDRSC without a network identifier specified can be located. In any case, the resource name is not passed to the alias name translation facility to perform name translation. However, VTAM can invoke the facility to determine if the owning SSCP is defined to the alias name translation facility.

With interconnected networks, defining CDRSCs becomes more complex as cross-network sessions, alias name translation, and adjacent SSCP rerouting are introduced. The configuration in [Figure 132 on page 450](#) shows two interconnected networks.




Figure 132. Example of two interconnected networks

To establish cross-network sessions between APPL1 in NETA and TERM1 in NETB, the following CDRM definitions are required if sessions are initiated from either HOST1 or HOST2. This CDRM major node is filed in HOST2 for the external CDRM HOST1.

HOST1	NETWORK	NETID=NETA
	CDRM	CDRDYN=YES, CDRSC=OPT
	GWPATH	GWN=GWNCP1, ADJNETSA=10, ADJNETEL=1, ...
	GWPATH	GWN=GWNCP2, ADJNETSA=10, ADJNETEL=1, ...

The following CDRM major node is filed in HOST1 for the external CDRM HOST2:

HOST2	NETWORK	NETID=NETB
	CDRM	CDRDYN=YES, CDRSC=OPT
	GWPATH	GWN=GWNCP1, ADJNETSA=20, ADJNETEL=1, ...
	GWPATH	GWN=GWNCP2, ADJNETSA=20, ADJNETEL=1, ...

Predefined cross-domain resources without network specification

If you do not code the NETWORK definition statement and you want to predefine CDRSCs (or model CDRSCs), you should specify the adjacent SSCP rather than the owning SSCP in the CDRM operand of the CDRSC definition statement.

Note: You cannot specify the NQNMODE or LUALIAS operand for cross-network resources without specifying a NETWORK definition statement.

The name can represent the alias or real name for the resource. When the real name and NETID are determined, VTAM updates the predefined CDRSC to represent the real name.

Note: If the real NETID is cross-network, no other same-network resource can be identified by the originally coded name.

When the last session for that CDRSC ends, VTAM restores the originally coded name so that the CDRSC is the same as that previously predefined. However, if NETID is coded, whether in this network or another, that indicates a real CDRSC.

You can define TERM1 (the destination LU in [Figure 131 on page 448](#)) in the following way:

	VBUILD	TYPE=CDRSC, ...
TERM1	CDRSC	CDRM=HOST21, ...

The CDRM, HOST21, is not the owning SSCP. Rather, it is the SSCP adjacent to the origin logical unit owning SSCP (HOST11). Because the CDRSC is defined in HOST11 network (or is allowed to default), the resource name (TERM1) is assumed to be the name known in HOST11 network (that is, the alias name).

Following are some advantages of predefining CDRSCs using this method:

- By coding the CDRM operand on the CDRSC definition statement, you can select the next SSCP in the session setup path. For simple networks that do not plan to provide alternate adjacent SSCPs, this method can eliminate the need for defining adjacent SSCP tables. For more complex configurations, this

method of defining CDRSCs identifies the first SSCP that should be tried when alternate adjacent SSCPs are available. In this case, VTAM reorders the adjacent SSCP table you defined; the first SSCP in the table is now the SSCP named on the CDRM operand of the CDRSC definition statement, causing this CDRM to be used first for routing.

- For some configurations, new CDRSC definitions are not required. If a multiple-network configuration is the result of splitting a single network and if the owning CDRMs have not changed, the CDRSC definitions previously defined still cause the request to be routed to the correct SSCP.

The disadvantage of predefining CDRSCs with this method is that the length of the session setup path is increased, though not as much as with dynamically defined CDRSCs using the resource alias name (because no definitions have to be built dynamically). When the real SSCP name and NETID are determined, VTAM updates the predefined CDRSC to represent the real name. When the last session for that CDRSC ends, VTAM restores the alias name so that the CDRSC is the same as originally predefined and restores the CDRM if it is coded. The NQNMODE and LUALIAS operands are only valid for cross-network CDRSCs with with an owning network specified.

This method of predefining CDRSCs is most suitable for simple configurations (for example, where one network is split into two) or for hosts that rely on another host for gateway support. In the example, HOST11 passes all session requests to HOST21 or HOST22. Those hosts are then responsible for alias name translation, alias address requests, and adjacent SSCP rerouting.

Predefined cross-domain resources with network specification

Predefining CDRSCs (or model CDRSCs) with network specification involves including NETWORK definition statements in the CDRSC definition to identify the network in which the resource is located. When a CDRSC definition follows a NETWORK definition statement, VTAM recognizes the specified CDRM as the actual owning SSCP and not as the next SSCP on the session setup path. A session with that CDRM is required only if it is also the next SSCP on the session setup path. The CDRSC definition statement can optionally include the name of the CDRM that owns the logical unit represented by the CDRSC.

In the following example, which is used to define TERM1 in HOST11, TERM1 is defined to be in NET3, and its owning CDRM is HOST31.

```
VBUILD TYPE=CDRSC,...  
NETWORK NETID=NET3  
TERM1   CDRSC CDRM=HOST31,...
```

VTAM recognizes that HOST31 is the actual owning SSCP and not the next SSCP on the session setup path.

Following are some advantages of using this method:

- In a configuration with no name conflicts, this method defines the shortest session setup path. However, unless you are using NQNMODE=NQNAME, a resource, even when defined with a NETID, cannot be predefined if its name is not uniquely known in the host network. If you are using NQNMODE=NQNAME, a resource, when defined with a NETID, can be predefined even if the name is duplicated in another network.

This method of predefining CDRSCs is designed for those systems that want to quickly locate a resource and use its real name. Therefore, the name must be uniquely known within the network in which it is defined. If the name duplicates a name in an interconnected network, it must be using its network-qualified name.

- By predefining CDRSCs in this way, you can define very specific adjacent SSCP tables. Because the NETID and, optionally, the destination SSCP are known, you can define adjacent SSCP tables for a specific NETID or NETID and SSCP name. This can eliminate the VTAM need for a default SSCP list, which might not provide the shortest session setup path.
- It adds security by restricting only authorized SSCPs to access resources in this host, assuming no dynamic CDRSCs are permitted.
- You can use NQNMODE=NQNAME to enable CDRSCs in different networks to have the same name.

Following are some disadvantages:

- If you specify the VFYOWNER=YES operand on the CDRSC definition statement, changing CDRMs in a backup situation can be complex. That is, if you have coded the CDRM operand on the CDRSC definition statement, VTAM assumes that it is the owning SSCP name. The operator must issue a MODIFY CDRM operator command in each network that contains a CDRSC definition, with the previous CDRM listed as the owner. Then (as in [Figure 131 on page 448](#)), if HOST31 fails and TERM1 is acquired by HOST32 but the CDRM is not updated in HOST11, the session setup attempt fails because HOST31 is still defined as the destination of the CDRM owning the logical unit. By specifying VFYOWNER=NO on the CDRSC definition statements along the route, VTAM does not verify if that CDRM is the owner but assumes that the session request contains the backup owner and thus routes the session request to that CDRM. Then, only the operators at the session origin need to modify the CDRSC owner to the backup CDRM owner.
- If you are not using NQNMODE=NQNAME, this method does not work for interconnected networks with duplicate names. The method requires that all the connected networks have unique resource names. In many cases, renaming and predefining CDRSCs using NETID or using network-qualified names is preferable to using the alias name translation facility.

This method is recommended for networks that contain unique names and require fast session setup.

If the CDRSC moves randomly among networks, specifying VFYOWNER=NO allows session setup to succeed, but added security is lost in that the intermediate VTAMs on the session setup path do not verify the owning host.

Dynamic cross-domain resources

Instead of predefining CDRSCs, you can let VTAM dynamically allocate and define some or all of them.

Dynamic CDRSC definition is used when VTAM receives a request to establish a session from or for an undefined cross-domain resource or for directory search requests. VTAM handles the request by dynamically creating a temporary CDRSC definition for the cross-domain resource. Without dynamic definition of CDRSCs, the number of resources requiring definition can be extremely large. If you provide networking services within an enterprise or for external users, this method is recommended.

For example, in [Figure 131 on page 448](#), if you choose not to predefine a CDRSC for TERM1 and the CDRM definition for HOST11 and HOST31 is set up to allow dynamic CDRSC definition, the following occurs:

1. A CDRSC for TERM1 is dynamically allocated in HOST11.
2. Unless you are using network-qualified names, if duplicate resource names exist, VTAM must select an alias name to see the session partner located in the other network. For example, if TERM1 exists in both NET1 and NET3, VTAM chooses an alias name whereby NET1 refers to the TERM1 located in NET3. VTAM uses the NetView alias name translation facility or the VTAM session management exit routine to identify the network in which the session partner resides. Then, VTAM resolves the alias name to the name used in the session partner network. Optionally, the NetView alias name translation facility can be used to obtain the name of the SSCP that owns the session partner.

If you are using network-qualified names, the NETID uniquely identifies the resources, and VTAM does not have to select an alias name for the session partner located in the other network.

3. VTAM uses an adjacent SSCP table to locate the next SSCP to which the session-setup request should be sent to reach the destination network.

When the session-initiation request leaves the first gateway (that is, the gateway NCP connecting NET1 and NET2), VTAM sends out a search to find the real name and NETID to replace the alias name for the destination logical unit if the alias name has not yet been translated into a real name and NETID. For information about how this affects routing, see [“Name assumption” on page 449](#).

If no translation takes place in the configuration in [Figure 131 on page 448](#), the session initiation request means that the destination network is assumed to be NET3 when the request is sent to either HOST31 or HOST32.

Following are some advantages of using this method:

- The number of CDRSC definition statements required in each host is greatly reduced. VTAM storage requirements are also reduced because only needed definitions exist.

- After the real name and network are known, the lack of predefined CDRSCs does not affect the length of the session setup path because the dynamic CDRSCs are retained for a user-specified time interval (CDRSCTI start option) after the last session terminates.

Following are the disadvantages of using this method:

- The main disadvantage of eliminating predefined CDRSCs is that the length of the session setup path for unqualified session requests is increased before VTAM identifies the destination logical unit real name and network. In this case, a CDRSC is dynamically allocated using the alias name. VTAM must assume that the name supplied to it in a session request is an alias name. After the real name is determined (by using the alias name translation facility or the alias function of the session management exit routine or from the session initiation response), the acquired information (the real name and network) is merged into the existing dynamically allocated CDRSC. Unless an alias name translation facility that supplies the name of the SSCP owning the session partner is used, each adjacent SSCP leading to the destination network is tried until the session is established or it is determined that none of the SSCPs owns that requested session partner. This process is repeated by each gateway-capable SSCP that receives the session initiation request.
- If an inactive LU is found which represents the destination LU, VTAM can allocate a dynamic CDRSC to enable the session initiation request to be sent to other VTAMs. The decision to allocate the dynamic CDRSC depends on several factors. If the session initiation request specified that the SSCP owning the destination LU was this host, the request fails. Also, if the LU is not eligible to be made into a shadow resource, the session is rejected.
- An adjacent SSCP table is required if you want to use dynamic CDRSC definition for destination LUs. You can predefine this table or specify the DYNASSCP start option to have VTAM create a dynamic table.

Using this method of dynamically allocating CDRSCs is most beneficial for intermediate and destination networks along the session setup path; the destination logical unit real name and network identifier are known by the time the session request reaches these hosts, so replacing CDRSCs is not required. Without dynamic definition of CDRSCs, the number of resources requiring definition can be extremely large. If you provide networking services within an enterprise or for external users, this method is recommended.

If the dynamic CDRSC definition is authorized, VTAM does the following actions:

- Builds a CDRSC minor node definition to represent the undefined LU
- Records the owning CDRM name (when determined)
- Initializes LU information from the session-initiation request or response

Dynamically defined CDRSC minor nodes are collected in a CDRSC major node named ISTCDRDY. ISTCDRDY is activated automatically during VTAM initialization, and deactivated automatically during VTAM termination. ISTCDRDY can be deactivated and activated by an operator command; this causes all currently active dynamic CDRSCs to be deleted. In addition, all sessions involving dynamically created CDRSCs (including CP-CP sessions with this host, if the partner CP was dynamically defined) are terminated.

While allowing dynamic CDRSC definition for session partners in another domain, you can still restrict access to application programs using the following exit routines:

- Session management exit routine (see [z/OS Communications Server: SNA Customization](#))
- Session authorization exit routine (see [z/OS Communications Server: SNA Customization](#))
- Application program logon exit routine (see [z/OS Communications Server: SNA Programming](#))

Dynamically defined CDRSCs are deactivated and deleted by VTAM on a periodic basis if they are not in use. That is, if a CDRSC has no active sessions and has had none for a defined interval of time, the definition is discarded. This interval is set by the CDRSCTI start option of VTAM.

A dynamically defined CDRSC that becomes a shadow resource has all of its sessions moved to the corresponding LU definition. Because of this, it is considered to have no active sessions as soon as it becomes a shadow resource and is then deleted.

If dynamically defined CDRSCs are used, initial session initiation time can be slightly longer unless network-qualified names are used in session initiation requests.

Defining adjacent SSCPs

This section contains information about adjacent SSCP tables that are used for cross-network session requests. For general information about adjacent SSCP tables, see [“Adjacent SSCPs” on page 431](#).

You can use adjacent SSCP tables in a multiple-network environment to:

- Reach destination SSCPs in other networks
- Determine the next SSCP in the session-initiation path to reach a destination network for an LU-LU session

Types of adjacent SSCP tables

There are three types of SSCP tables that are used for cross-network session requests:

Network-specific list

The network-specific list has a NETWORK definition statement coded in the adjacent SSCP table, and the NETID operand specifies the NETID of the destination network.

Network-specific and CDRM-specific list

The network-specific and CDRM-specific list specifies both a destination network and a destination CDRM.

Default list

The default SSCP list is used whenever one of the following conditions is true:

- The destination network is not known.
- The destination network is known, but there is not a network-specific table associated with the network.
- The destination network and owning SSCP name are known, but the specific table cannot be located.

When a CDRSC network is known, but its owner is not, VTAM uses the network-specific adjacent SSCP table for that network. If none is defined, the default is used.

The default SSCP list is part of the adjacent SSCP table. You can include network-specific SSCP lists for each network and a default list for use when the real network is not known.

Deciding whether to code adjacent SSCP tables

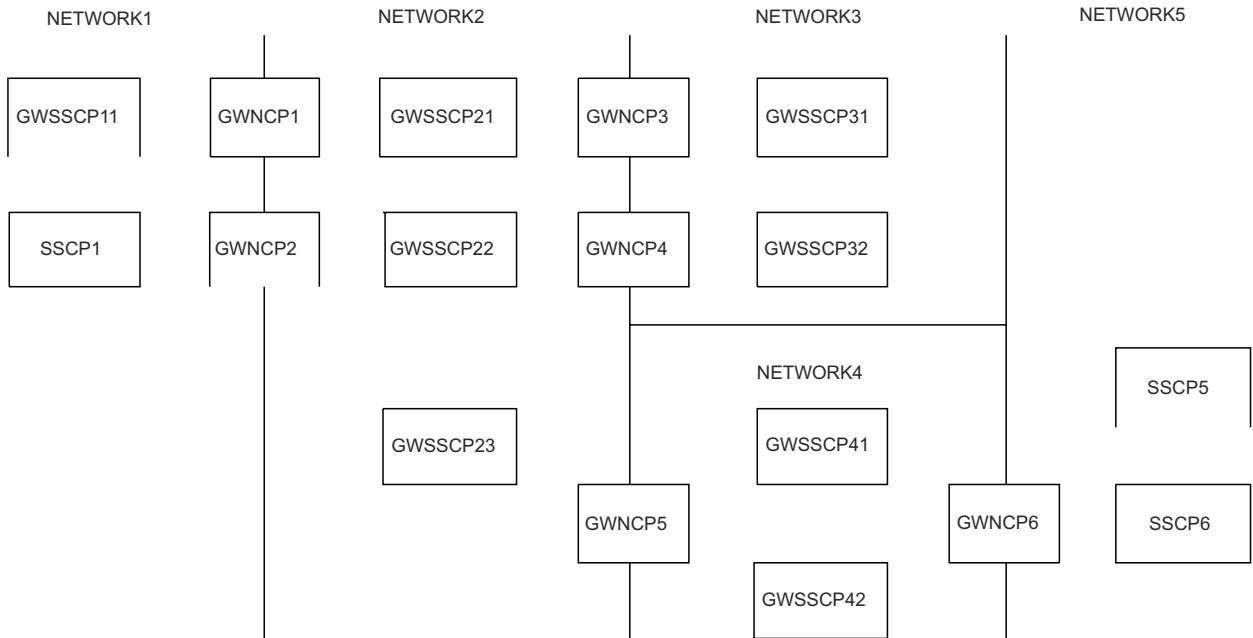
A host does not need an adjacent SSCP list for a network if the following are true:

- The host SSCP is directly in session with the destination SSCP (that is, there are no intermediate SSCPs on the session path), and the host SSCP knows the destination SSCP name.
- Predefined CDRSCs without NETWORK, but with CDRM, are specified.
- An alias name translation program or the alias function of the session management exit routine is used to find the destination SSCP name.

Note: You might need to define adjacent SSCP tables containing the type 2.1 network ID for session requests destined for a nonnative network logical unit.

Sample of adjacent SSCP tables for a multiple-network environment

The sample adjacent SSCP table definitions included below are written for the configuration in [Figure 133 on page 455](#).



Note: The connections between hosts and NCPs have been omitted for clarity. You can assume that each NCP connects to each host in the networks it joins.

Figure 133. Multiple-network configuration

The adjacent SSCP table in the following example applies to GWSSCP21 in NETWORK2 and contains default SSCP lists for all the networks shown in [Figure 133 on page 455](#).

```
TABLE2  VBUILD  TYPE=ADJSSCP
*****
*  DEFAULT SSCP LIST: THIS LIST IS USED WHEN THE CDRM          *
*                      IS UNKNOWN AND NETWORK IS UNKNOWN      *
*****
      NETWORK          UNNAMED NETWORK DEFINITION
GWSSCP11 ADJCDRM      WITH DEFAULT ADJACENT SSCP
GWSSCP31 ADJCDRM      LIST FOR ALL NETWORKS
GWSSCP32 ADJCDRM
GWSSCP41 ADJCDRM
GWSSCP42 ADJCDRM
*
*****
*  DEFAULT SSCP LIST FOR NETWORK1: USED WHEN THE CDRM IS UNKNOWN *
*                      AND NETWORK IS KNOWN TO BE NETWORK1      *
*****
      NETWORK  NETID=NETWORK1  DESTINATION NETWORK
GWSSCP11 ADJCDRM              ADJACENT SSCP FOR NETWORK1
*
*****
*  DEFAULT SSCP LIST FOR NETWORK2: USED WHEN THE CDRM IS UNKNOWN *
*                      AND NETWORK IS KNOWN TO BE NETWORK2      *
*****
      NETWORK  NETID=NETWORK2  DESTINATION NETWORK
GWSSCP22 ADJCDRM              DEFAULT ADJACENT SSCP LIST
GWSSCP23 ADJCDRM              FOR NETWORK2
*
*****
*  DEFAULT SSCP LIST FOR NETWORK3: USED WHEN THE CDRM IS UNKNOWN *
*                      AND NETWORK IS KNOWN TO BE NETWORK3      *
*****
      NETWORK  NETID=NETWORK3  DESTINATION NETWORK
GWSSCP32 ADJCDRM              DEFAULT ADJACENT SSCP LIST
GWSSCP31 ADJCDRM              FOR NETWORK3
*
*****
*  DEFAULT SSCP LIST FOR NETWORK4 OR NETWORK5: CDRM IS UNKNOWN AND *
*                      NETWORK IS KNOWN TO BE NETWORK4 OR 5*
*****
      NETWORK  NETID=NETWORK4  DESTINATION NETWORK
      NETWORK  NETID=NETWORK5  DESTINATION NETWORK
```

```

GWSSCP41  ADJCDRM                DEFAULT ADJACENT SSCP LIST
GWSSCP42  ADJCDRM                FOR NETWORK4 AND NETWORK5
*
*****
*   CDRM-SPECIFIC LIST FOR SSCP5: CDRM IS KNOWN                *
*   *                   TO BE SSCP5 IN NETWORK5                *
*****
SSCP5      CDRM                DESTINATION SSCP
GWSSCP42  ADJCDRM                ADJACENT SSCP LIST
GWSSCP41  ADJCDRM                FOR SSCP5 IN NETWORK5

```

Example default lists

In the above adjacent SSCP table, the default SSCP list for all networks is:

```

GWSSCP11
GWSSCP31
GWSSCP32
GWSSCP41
GWSSCP42

```

This list is used whenever the following are true:

- The destination logical unit network is not known.
- The destination network ID is known, but no entries correspond to the destination network.

The default SSCP list for NETWORK4 and NETWORK5 is as follows:

```

GWSSCP41
GWSSCP42

```

This list is used whenever VTAM knows the destination network to be either NETWORK4 or NETWORK5, but does not know the destination SSCP name.

If the destination SSCP is known to be SSCP5 in NETWORK5, the following adjacent SSCP list is used instead:

```

GWSSCP42
GWSSCP41

```

Example network-specific lists

The rest of the adjacent SSCP table is used when VTAM sets up sessions with CDRSCs whose owning CDRM and network are known. For example, if GWSSCP21 has the following definition statement coded in its CDRSC major node:

```

TS002      NETWORK  NETID=NETWORK3
           CDRSC    CDRM=GWSSCP32

```

Because the CDRM is coded on the CDRSC definition statement, VTAM assumes that this CDRM is the true owner of the resource. Because the NETWORK definition statement is coded, VTAM uses the coded CDRM and the network-specific SSCP table to try to locate the CDRSC.

VTAM first builds an adjacent SSCP table for the CDRSC. The CDRM is defined as SSCP1, so GWSSCP32 is the first entry in the list. The rest of the list consists of the SSCPs that are adjacent to GWSSCP32, as defined in the default portion of the sample table:

```

GWSSCP32  CDRM
GWSSCP31  ADJCDRM

```

If the session setup fails through GWSSCP32, VTAM routes the request through GWSSCP31.

If no network had been coded for TSO02, VTAM would have used the default SSCP list (GWSSCP11, GWSSCP31, GWSSCP32, GWSSCP41, GWSSCP42). This could result in some session setup requests being routed outside the network through the gateway SSCPs. The adjacent SSCP table can thus prevent

time and resources from being used in a search that might be unnecessary. The session management SSCP selection function can also be used to alter the SSCP list.

For VTAM Version 3 Release 2 and later releases, an ADJSSCP table defined with a NETWORK definition statement for the host network is used only when the destination resource is known to reside in the host network. If a destination resource is known to reside in NETWORK2, the list of ADJSSCPs used for a session request is:

GWSSCP22
GWSSCP23

Before VTAM Version 3 Release 2, an ADJSSCP table defined with a NETWORK definition statement for the host network is treated as a default table.

Overriding the SSCP lists

If a list of adjacent SSCP exists for a destination network, an SSCP, or both, and if VTAM already has an SSCP-SSCP session with the destination SSCP, session setup processing can override the list; that is, the destination SSCP becomes the first choice for the specific session request. This override is temporary and does not change the actual adjacent SSCP table that you define.

Thus, the default SSCP list (TABLE2) includes the following definition statements:

NETWORK	NETID=NETWORK3	DESTINATION NETWORK
GWSSCP32 ADJCDRM		DEFAULT ADJACENT SSCP LIST
GWSSCP31 ADJCDRM		FOR NETWORK3

Assume that a cross-network SSCP-SSCP session exists between GWSSCP21 and GWSSCP31. If GWSSCP21 receives a cross-network initiate request for a session with a logical unit owned by GWSSCP31 in NETWORK3, the list of adjacent SSCP in the table (GWSSCP32 and GWSSCP31) is reordered for this session setup as follows:

GWSSCP31
GWSSCP32

The actual table you define is unaltered.

The destination SSCP name, GWSSCP31, is available in the initiate request if the alias name translation facility or a predefined CDRSC with CDRM name is used.

Request routing

Session initiation and INQUIRE APPSTAT requests are routed from the originating LU SSCP to the destination LU SSCP using trial-and-error routing. There are several different means of determining the list of adjacent SSCP to try when performing trial-and-error routing. However, after the list is determined, the processing is very similar for each of these RUs.

Rerouting of requests

If the adjacent SSCP to which a cross-network session setup request is routed does not own the destination logical unit, it reroutes the request, provided that the following are true:

- An appropriate adjacent SSCP is available.
- The maximum SSCP rerouting count has not been reached.
- This SSCP is gateway capable.

Otherwise, the request is rejected.

If the SSCP that sent the session setup request is in the list, the entry is not used to reroute the request. This prevents a request from looping. The actual table you define is unaltered.

Rules for cross-network session request routing

The rules for routing cross-network session requests are as follows:

- The request is not routed if its visit count (SSCP rerouting count) has reached 0.
- The request is never routed back to the SSCP from which it was received.
- The request is never routed back to its originating SSCP.
- The request is never routed in this network if it was received from an SSCP in this network.
- An SSCP receiving a session request from another network can route the request to a different SSCP in that network through a different gateway NCP. It cannot route the request back using the same gateway NCP.
- The request received by a nongateway SSCP cannot be rerouted by that SSCP.
- The cross-network request received by a gateway-capable SSCP can be rerouted by that SSCP.
- A nongateway SSCP can perform trial-and-error routing for a session that it initiated.
- If the real NETID of the DLU is known, routing is not attempted to a nongateway SSCP with a different NETID, unless the adjacent SSCP supports the nonnative network connection function.
- A request is not routed to an SSCP that does not support the required function. (For example, automatic logon sessions are not routed to SSCPs that do not support automatic logon sessions.)

Dynamically defined CDRSCs and adjacent SSCP tables

If dynamic definition of resources is allowed and if no predefined CDRSC is located, VTAM automatically creates a CDRSC for the logical unit as part of routing. After the owning SSCP is found by trial-and-error routing of the session initiation request, the dynamically defined CDRSC is updated with the real network ID and owning SSCP name.

For releases of VTAM before Version 4, if two hosts attempt to establish a cross-network session with a logical unit that is not predefined in either of the initiating hosts, a conflict can arise in any intermediate host that receives these requests. The receiving host will then issue a sense code indicating duplicate DLU names. The sessions can be successfully established if the requests can be rerouted. This situation can be avoided, however, by predefining the DLU with NETID or by using the NetView alias name translation facility or the alias function of the session management exit routine to define the name of the resource instead of allowing it to be dynamically defined.

Alias name translation and adjacent SSCP tables

The NetView alias name translation facility can be used to determine the owning SSCP for a requested resource in the same network. To use this facility for an LU, define the LU and its owning SSCP to the alias name translation facility, because the alias name translation facility cannot translate a name as known in a network to a different name in the same network.

Define each alias name to be the same as the real name. During a session setup request, if VTAM does not have a definition of a destination logical unit or does not know the owning SSCP for the resource, VTAM can invoke the alias name translation facility with the name of the logical unit. The alias name translation facility returns the name of the SSCP defined as the owner of the logical unit, and the destination network identifier and the logical unit alias name in that network. Because this system has only one network, the network identifier is the VTAM network identifier, and the alias name is the same as the real name. Thus, the alias name translation facility is used as a directory to locate the owning SSCP and can direct the session setup request to the proper VTAM.

Cross-network routing

In a multiple-network environment, application programs and terminals managed in one network can access SNA resources controlled and managed in another network, but each network can preserve its autonomy and maintain unique network characteristics, such as the network address structure and network naming conventions.

Network address structures

When SNA networks are interconnected, each can maintain its current network addressing scheme. However, there are some addressing considerations.

Maintaining address structures

The network addresses assigned to nodes (host processors or communication controllers) in one network can be duplicated in other connected networks, as can the following subarea addressing capabilities:

- Extended subarea addressing
- Extended network addressing
- Nonextended network addressing (MAXSUBA)

The maximum number of addressable network nodes or subareas does not have to be consistent or compatible across the interconnected networks.

SNA network interconnection (SNI) eliminates the architectural requirement that interconnected networks must use a common network address structure with all unique network addresses. Instead, it permits each network to establish its own network addressing capabilities.

Addressing considerations for interconnected networks

Any network that has greater addressing capability than its adjacent networks must have extended subarea addressing-capable gateway NCPs and gateway SSCPs that are involved in setting up cross-network sessions for high subareas. Specifically, consider the following situations:

- The gateway NCP addressing capability in each network must be at least equal to the largest subarea address in the network, and its subarea address in the network must be within the range supported by the subarea with the lowest addressing capabilities. This restriction applies only for subareas that are using the gateway NCP for cross-network sessions.
- The gateway VTAM addressing capability must be greater than or equal to the largest subarea address in the network, and its subarea address must be within the range supported by the SSCP with the lowest addressing capabilities. That is, the gateway SSCP must be capable of supporting a CDRM session with any SSCP in the network, and it must be capable of owning the gateway NCPs it is to control.
- To properly support communication between networks where either network contains subareas that support cross-network sessions and that have subarea addresses above the capabilities of any of the other subareas involved in supporting or setting up the cross-network session, the gateway SSCPs in each network must be capable of supporting extended subarea addressing.
- For an extended subarea addressing gateway NCP to successfully send a CONTACT request unit to an adjacent cross-network gateway NCP with a subarea address greater than 255, the link station must be owned by an extended subarea addressing-capable VTAM.

Network naming conventions

SNA requires each terminal and application program (logical unit) to be identified by a unique, network-qualified name. Network-qualified names are used by specifying the NQNMODE=NQNAME start option or defining cross-network resources with the NQNMODE=NQNAME operand on the CDRSC definition statement.

If you choose not to use network-qualified names, you can use the optional facility SNI provides for translating terminal and application program names to cross-network (alias) names to avoid name duplication conflicts. By using this facility, the name structure used in each network can be maintained. A common network naming convention among the interconnected networks does not have to be established. If multiple SNA resources are known by the same name in different networks, the name conflict can be resolved by using one or a combination of:

- The NetView alias name translation facility
- The VTAM session management exit routine

- Network-qualified names

If you choose to fully implement network-qualified names, you do not need to use either the NetView name translation facility or the VTAM session management exit routine to resolve duplicate names.

Controlling paths for interconnected networks

For SNA network interconnection (SNI), paths are defined the same as for single- and multiple-domain environments within each network (see [“Network routing for subarea nodes”](#) on page 256 and [“How to plan routes in your network”](#) on page 264). Rules for explicit and virtual routes do not cross network boundaries. For example, in [Figure 134 on page 460](#) assume that you are setting up a session between LU1 and LU2 and that they are in the same network. The same explicit route that is used from HOST1 to GWNCBPB (ER0) would have to be used from GWNCBPB to HOST2. However, because HOST1 and HOST2 are in different networks, the two explicit route definitions are independent of each other.

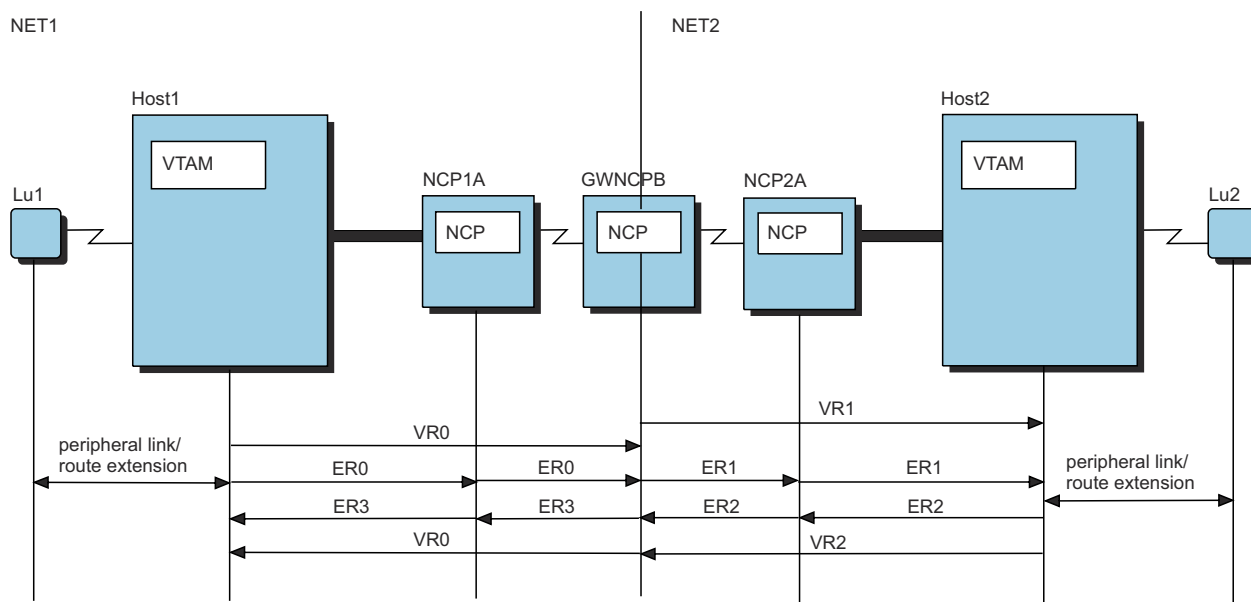


Figure 134. Multiple-network paths

Paths do not cross network boundaries. Cross-network routing requires separate paths in each network, each having the gateway NCP as the endpoint. If you are combining networks, you might have to add paths in nongateway NCPs that go to the gateway NCP.

If you are splitting an existing network into two networks, redefine the paths so that each path passing through a gateway NCP is defined separately in each of the networks.

Gateway path selection

For cross-network sessions, the gateway path selection function of the session management exit routine can be used to modify a list of alternate gateway paths to distribute LU-LU sessions among two or more gateway NCPs. You can shorten the list by deleting entries or reorder it by putting the preferred paths at the top of the list.

If you do not provide this function in your session management exit routine, VTAM uses the list of gateway paths determined by the GWPATH definition statements. However, the gateway NCP used for the SSCP-SSCP session is always the first one used for LU-LU sessions. For information about using the session management exit routine, see [z/OS Communications Server: SNA Customization](#).

A gateway path is selected by one of the following methods:

SSCP-SSCP sessions

Each successive gateway path between two VTAMs is tried until either the SSCP-SSCP session is established or all of the possible session paths have been tried unsuccessfully.

The class of service name that is normally used is the COS name for SSCP sessions, ISTVTCOS. This is the default name that is used for establishing a cross-network SSCP-SSCP session and, if used, is the COS name that is passed to the NetView alias name translation facility or the VTAM session management exit routine. You can specify a different Class of Service name by coding the value on the ADJNETCS operand on the GWPATH definition statement. VTAM uses that COS name to locate the virtual route list in the proper COS table as specified on the COSTAB operand of the BUILD or NETWORK definition statements in NCP. If the ADJNETCS operand is specified, VTAM does not invoke the NetView alias name translation facility or the VTAM session management exit routine. For more information about network routing in a multiple-network environment, see [“Cross-network routing” on page 458](#).

LU-LU sessions

Each gateway NCP (defined by a GWPATH definition statement) between the two VTAMs is tried until one of the following occurs:

- The requesting logical unit receives a positive response to the initiate request for the session, meaning that the requested session partner has been located.
- All the possible gateway NCPs (defined by the GWPATH definition statements for the CDRM in another network) have been tried unsuccessfully.
- A failure not related to the GWNCP is detected.

If multiple gateway paths are used, VTAM selects a particular gateway NCP for each LU-LU session. The gateway NCP used for the cross-network SSCP-SSCP session is always the one VTAM attempts first for LU-LU session setup. If this is unsuccessful, the gateway VTAM selects another gateway NCP for the session. If you coded a session management exit routine to shorten or reorder the list of gateway NCPs, the gateway VTAM uses that altered list to establish the LU-LU session path.

During session establishment of a cross-network LU-LU session, the gateway VTAM that is controlling gateway path selection gets the list of alternate gateway NCPs defined by the GWPATH definition statements. If the BIND RU that flows between the requesting logical unit and the session partner fails because a virtual route required by the session cannot be activated, no further gateway path selection is attempted. VTAM messages are issued to inform the operator that session establishment failed, and the requesting logical unit is informed.

Coding the gateway path

The GWPATH definition statement is required for all gateway VTAMs that initiate cross-network SSCP-SSCP sessions. If you do not code GWPATH definition statements following a CDRM definition statement for a VTAM in another network, code the SUBAREA operand on that CDRM definition statement so that VTAM can determine the gateway NCP to use.

Dynamic cross-domain resource definition must be allowed for the host CDRM if the NetView alias name translation facility or the alias function in the session management exit is used to translate logical unit names.

Handling class of service tables

Following are several important issues to consider when using COS tables for interconnected networks:

- Using COS tables for gateway VTAMs and gateway NCPs
- Using multiple identical COS tables (that is, using one COS table for more than one network or for more than one subarea node within a network)
- Using different COS tables for interconnected networks
- Using conflicting COS table names

You can define different COS tables for VTAM to use within its own network or within gateway NCPs that it owns. When the PLU is in the VTAM subarea, VTAM uses ISTSDCOS to resolve the Class of Service name to a virtual route list. The virtual routes originating in the host subarea extend to the SLU subarea node or to a gateway NCP providing a path to the SLU.

When VTAM is acting as a gateway VTAM that has been designated to resolve COS names for a gateway NCP, VTAM uses the COS table named on the NETWORK or BUILD definition statement in the gateway NCP generation definition (whichever applies to the adjacent network to which you are routing). The gateway VTAM passes the virtual route list to the gateway NCP, and the gateway NCP performs the route activation for one of the routes in the resulting virtual route list within that network.

To establish a cross-network SSCP-SSCP session, VTAM uses the ISTVTCOS Class of Service name to establish a route to the adjacent network SSCP. On the GWPATH definition statement that identifies the gateway NCP to be used for the session, you can specify the ADJNETCS operand to identify a COS name other than ISTVTCOS. VTAM uses this COS name to determine the specific route in the adjacent network for the session.

All of the COS tables to be used for gateway NCPs controlled by a gateway VTAM must be available to that SSCP. If VTAM attempts to load a COS table that it cannot find, session requests associated with the table fail.

If more than one gateway VTAM shares control of a gateway NCP, those gateway VTAMs can be defined in such a way that some of them allow the others to resolve the COS name. This is specified by the GWCTL operand on the PCCU definition statement. In this case, only those gateway VTAMs responsible for resolving COS names for that gateway NCP must have the COS tables stored in their hosts.

For example, in Figure 135 on page 462, if Host1 has GWCTL=ONLY coded on the PCCU definition statement for GWNCPB, Host1 COS table is used for an SSCP-SSCP session between Host1 and Host2. If Host1 and Host2 have GWCTL=SHR coded on the PCCU definition statements for GWNCPB, the originating VTAM COS table is used.

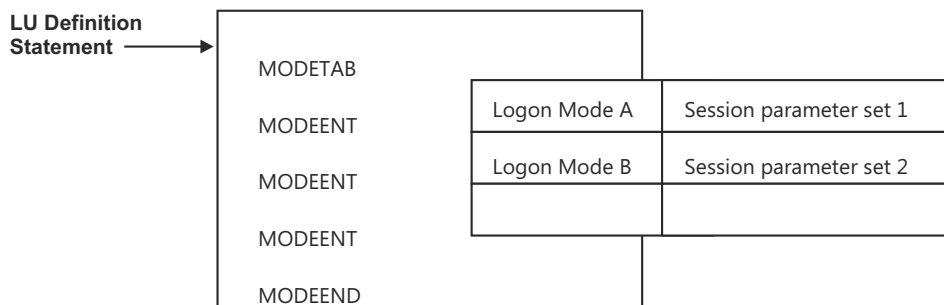


Figure 135. COS resolution in a multiple-network environment

The same COS table can be used for more than one network. You can also create a universal COS table to be used for all networks. This is possible if the COS names are the same and if virtual routes originating in different subarea nodes have identical VR numbers for routes providing the same levels of service.

If the same COS table can be used in more than one network connected to a gateway NCP, code the same table name on the COSTAB operands for the BUILD or NETWORK definition statements representing those networks in the NCP major node. VTAM loads only a single copy of the table into the VTAM storage.

If your interconnected networks require different COS tables, you can associate these table names with their proper networks by coding the table names on the COSTAB operands of the appropriate BUILD or NETWORK definition statements in the NCP generation definition.

The COS table for routes originating in a VTAM host must be named ISTSDCOS. If the same COS table can be used for a gateway NCP that is controlled by a different VTAM host, the COS table must be stored in the other VTAM host under a name other than ISTSDCOS (unless that VTAM can also share the same COS table for routes originating in that host).

In Figure 136 on page 463, networks NETA and NETB are connected in a back-to-back configuration.

The session between the PLU in NETA and the SLU in NETB uses the following three routes:

- A route within NETA from the PLU subarea to the gateway NCP (GWN1) subarea in NETA
- A route from the NETX subarea within GWN1 to the NETX subarea within GWN2
- A route from the NETB subarea within GWN2 to the SLU subarea




Figure 136. COS tables and routing in an SNI back-to-back configuration

Assuming that GWN1 resources reside in NETA (that is, NETA is GWN1 native network), GWN1 generation definition contains the following statements:

```
BUILD    NETID=NETA, ..., COSTAB=COSNETA1
...
NETWORK  NETID=NETX, ..., COSTAB=COSNETX1
```

Also, assuming that GWN2 native network is NETB, GWN2 generation definition contains the following statements:

```
BUILD    NETID=NETB, ..., COSTAB=COSNETB2
...
NETWORK  NETID=NETX, ..., COSTAB=COSNETX2
```

The following COS tables are used:

- In GWSSCPA, the SSCP of the PLU uses ISTSDCOS to select the list of virtual routes in NETA from which a route is selected for the session.
- GWSSCPA, as the gateway VTAM controlling GWN1, resolves the COS name for NETX to a list of virtual routes in NETX. It uses the table COSNETX1.
- GWSSCPB, as the gateway VTAM controlling GWN2, resolves the COS name for NETB to a list of virtual routes in NETB. It uses the table COSNETB2.

VTAM uses its own default virtual route list when a cross-network SSCP-SSCP session activation is attempted with a COS name other than ISTVTCOS or the unnamed Class of Service. Session activation proceeds, but a message is sent to warn the operator that the COS table should be corrected. Whenever convenient, you should correct the problem by restarting VTAM.

Note: For user session requests, you can specify a substitute set of COS parameters that are used if the COS name is unknown to this VTAM. See [“How session traffic is assigned to a specific route”](#) on page 260.

Address translation

A gateway NCP uses one subarea for each network that is to be connected. Thus, a gateway NCP appears to reside in more than one subarea. Associated with that subarea is a pool of element addresses that can be used for cross-network sessions. When a cross-network session is requested, the gateway NCP assigns a pair of element addresses, one in each network on either side of the gateway NCP, to the session partners. This enables the gateway NCP to create within each network a representation of the session partner in the other network.

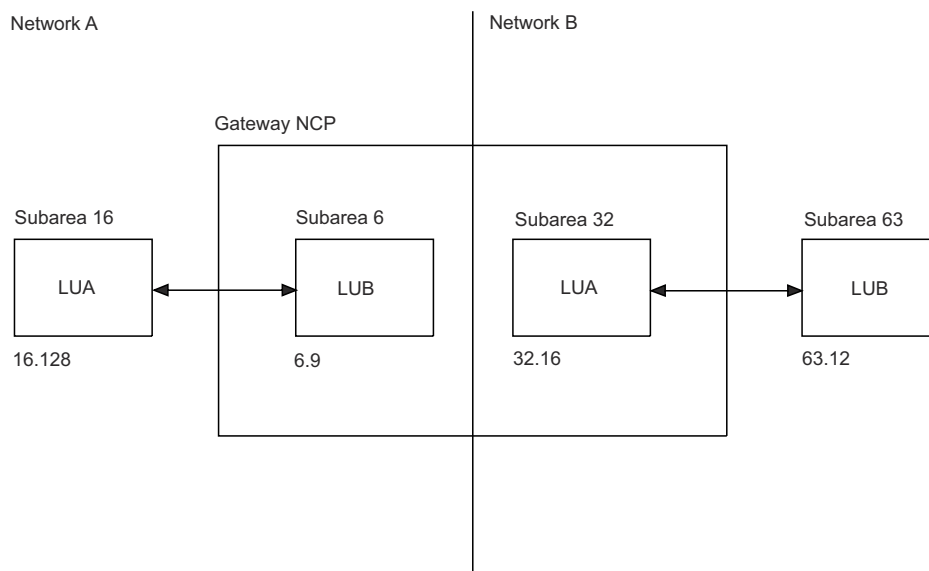


Figure 137. Address translation

For example, [Figure 137 on page 464](#) shows two networks (networks A and B) that are connected through one gateway NCP. An LU in Network A (LUA) is in session with another LU in Network B (LUB). LUB has a subarea address of 63 and an element address of 12 in its own network. Because these addresses might not be valid in Network A, the gateway NCP has given LUB a subarea number of 6 (the subarea of the gateway NCP in Network A) and an element address of 9. As a result, LUB now has a valid address in Network A that can be used for routing session data. When an LU in Network A sends data to LUB, it is routed to subarea 6, element 9, in the gateway NCP. The gateway NCP translates this address into the address to be used for LUB in Network B. It then passes the session data over to Network B. The session data is then routed through Network B to LUB in the usual manner.

Similarly, the gateway NCP has given LUA an address in Network B. Any session data destined for LUA from Network B is sent to this address.

Element addresses

Address translation is also done for cross-network SSCP-SSCP sessions. Each LU and SSCP with which you want to have a cross-network session requires at least one element address. The number of addresses that can be used to represent cross-network resources within a given network depends on the following situations:

- If extended network addressing is used, 32 768 element addresses can be used to assign alias addresses.
- If nonextended network addressing is used, the element pool is determined by the MAXSUBA specification. For example, if a network has MAXSUBA=63, each subarea can have at most 1024 element addresses.

Resources attached to the gateway NCP require element addresses in the native network. Therefore, while the number of possible cross-network sessions is not equal to the maximum number of elements in each subarea, the two must be considered together.

Parallel sessions and element addressing

Application programs that use nonparallel sessions require only one address to represent the primary LU to all of the secondary LUs. However, application programs that use parallel sessions require a unique address to represent the primary LU to each one of the secondary LUs. If you plan to have cross-network parallel sessions, the network containing the secondary LU must have one element address for each parallel session with the primary LU. In the network of the primary LU, only one element address is required for all sessions with the secondary LU. Therefore, the number of addresses you should reserve

in any particular network to represent application programs in other networks depends on whether the application program requires a unique primary LU address for each of its sessions.

For example, if terminals in your network are logged on to an application program such as CICS, you need only one element address in the gateway NCP to represent that CICS to your network. If parallel sessions are established between your CICS system and the CICS in the other network, the PLU must use an address for each session. If your CICS is acting as the SLU for these parallel sessions, an element address is required for each session in the gateway NCP to represent the PLU to your network. If the CICS in the other network is acting as the SLU for these parallel sessions, only one element address is required in the gateway NCP to represent that CICS as the SLU to your network.

The gateway NCP needs an additional address for the other network to represent each of your CICS parallel sessions to the other network. The gateway NCP also needs element addresses to represent your terminals to the network containing the other CICS.

Multiple ACBs and element addressing

If terminals in your network are logged on to an application program that opens a separate ACB for each user (such as the NetView or TSO programs), the number of element addresses that you need to represent the NetView program in your network is equal to the maximum number of terminals in your network that are concurrently logged on to that application program. In addition, the network containing the application program must have an element address for each of your terminals.

Resource name translation

SNA requires that each network resource be identified by a unique name. This requirement also applies to a multiple-network environment, unless you are using network-qualified names. Network-qualified names are implemented by specifying the NQNMODE=NQNAME start option or defining cross-network resources with the NQNMODE=NQNAME operand on the CDRSC definition statement. If you do not use network-qualified names, cross-network resources cannot be identified by the same name. For example, a problem occurs when a logical unit in one network requires access to a resource in another network that uses the same name to identify a different logical unit.

The class of service (COS) name and the logon mode entry name are also important when establishing sessions. The same COS and logon mode entry names must be used in each of the interconnected networks to establish cross-network sessions.

To meet these naming requirements when you are connecting networks, do one of the following steps:

- Redefine any duplicate resource names to preserve name uniqueness, and ensure that you use the same COS and logon mode entry names as the attached networks. (If you are using network-qualified names, you do not need to redefine resources to resolve duplicate names.)
- Define any resources with duplicate names so that no more than one of them is specified (either by definition or by start option) with NQNMODE=NAME and the others are specified (either by definition or by start option) with NQNMODE=NQNAME.
- Use the VTAM alias selection function in the session management exit routine to translate any duplicate resource names to unique resource names, and to map unlike COS and logon mode entry names in one network to the proper equivalents in the other network. For details about the session management exit, see [z/OS Communications Server: SNA Customization](#). (If you are using network-qualified names, you do not need to use the session management exit routine to resolve duplicate names.)
- Use the NetView alias name translation facility to translate any duplicate resource names to unique resource names and map unlike COS and logon mode entry names in one network to the proper equivalents in the other network. (If you are using network-qualified names, you do not need to use the NetView alias name translation facility to resolve duplicate names.)

The resource names that can be translated by the NetView alias name translation facility or in the session management exit routine are:

- Logical unit names
- Logon mode names

- Class of Service names

Alias selection function of the session management exit routine

You can use the alias function of the session management exit routine to perform name translation. By using this exit routine, you can provide your own alias and real names for logical units, logon mode names, Class of Service names, and SSCP names, instead of using the NetView alias name translation definitions. You can also indicate whether the NetView alias name translation facility is to be invoked for additional translation.

The session management exit routine can be invoked to translate a real or alias logical unit name associated with one network to a real or alias logical unit name for the destination network. You can map different LOGMODE and COS names in different networks to be functionally equivalent, and can also provide directory services by providing the owning SSCP names for the destination resource.

If you are using the NetView alias name translation facility, you can selectively invoke it to perform additional translations not performed in the exit routine. When the exit is invoked, there is an indicator that informs you if the NetView alias name translation facility is currently active. The exit routine can determine whether the session requests should be rejected, depending on the need for further translation by the NetView alias name translation facility and whether it is active. It can also accept or reject the use of a substitute COS.

For information about using the alias selection function of the session management exit routine, see [z/OS Communications Server: SNA Customization](#).

NetView alias name translation facility

The NetView program provides the alias name translation facility that can:

- Translate any LU name duplication conflicts
- Map logon mode and COS names

Translating LU names

No two LUs in an SNA network can have the same name. When you connect independent SNA networks, however, there might be some duplicate names among these networks. To provide the unique LU names required for SNA sessions but still allow each network to have independent resource names, interconnection provides for name translation between networks. Logical unit names are translated from names known in one network to names known in the network of the session partner to avoid duplicate names.

Network-qualified names provide an option to name translation and enable you to have resources by the same name in different networks without requiring you to use name translation.

If you are not using network-qualified names, you can use name translation, which allows each LU to have two names: a real name by which it is known in its own network and an alias name by which it is known in another network. For example, an LU in NETA is known by its real name LUA1A3 in NETA, and an LU in NETB is also known by its real name LUA1A3 in NETB. If NETA is connected to NETB, a name duplication conflict develops. The alias name translation facility enables LUA1A3 in NETA to be identified by an alias name in NETB (for example, BBA1A3). By using this procedure, you can maintain name uniqueness within a network and adhere to SNA naming requirements.

The NetView alias name translation facility uses installation-defined translation tables. The tables are a one-to-one mapping of the alias LU name to the real LU name and their corresponding networks. This technique differs somewhat from that used for maintaining alias addresses across a network boundary. That is, an exact one-to-one mapping is defined by an installation using the translation tables. However, terminal users and network operators do not normally know an LU by its network subarea and element address. Instead, it is known by its unique name in the network. If the alias name (by which an LU is known) is able to change dynamically with each session, network operations might be very difficult, and network security might be affected. The NetView translation tables therefore predefine the relationship between an alias name and a real name.

Note: The alias name translation facility cannot translate a name as known in one network to a different name in the same network. Also, because NetView domain identifiers cannot be translated, they must be unique for all NetView programs that have cross-network sessions. Resource names for LU 6.2 sessions also cannot be translated.

Mapping logon mode names

Before a session is established, both session partners must agree on a common set of session parameters. They do this by specifying a logon mode name that identifies a particular set of session parameters. Within an SNA network, both session partners use the same logon mode name, thus ensuring that the correct session parameters are used. When you connect independent SNA networks, however, a given set of session parameters might be known by different names in different networks. Also, a given name might see different session parameters in each network. The alias name translation facility can be used to map a logon mode name in one network to another name in the destination network that provides equivalent session parameters. This translation can occur in any SSCP on the setup path, after the network identifier of the destination network is determined.

As usual, the logon mode name that is used to establish a session must exist in the logon mode table of the secondary logical unit SSCP.

Mapping Class of Service names

Before a session is established, both session partners must agree on a common Class of Service that defines a particular ordered list of virtual routes. They do this by specifying a Class of Service name. Within an SNA network, both session partners can use the same Class of Service name, thus ensuring that the correct Class of Service is used. When you connect independent SNA networks, however, a given Class of Service might be known by different names in different networks, or the same name might be used for different Classes of Service. When the alias name translation facility is used, a Class of Service name in one network can be mapped to a different name in each adjacent network along a cross-network session path to provide equivalent service. This translation can occur in any SSCP on the setup path. For SLU initiated sessions, you must also map the COS name for the origin logical unit (OLU) to the destination logical unit (DLU).

The COS name entry used for the session must exist in the COS table of the primary logical unit SSCP and in the gateway VTAM that performs the SETCV. For information about determining which gateway VTAM performs the SETCV, see [“Cross-network routing” on page 458](#).

The NetView alias name translation facility can be used to establish equivalence for the required session parameters. The NetView alias name translation facility maps the logon mode name in one network to the logon mode name used in the other network. The same process is used for the COS name. If the COS name needed to establish a virtual route and transmission priority in one network does not exist in the COS table of the PLU in the other network, the NetView alias name translation facility can be used.

This alias name translation facility maps the COS name in one network to the COS name in the other network that provides a network route with equivalent characteristics (for example, high speed, high availability, and security).

A substitute COS entry can also be used. The substitute COS entry, chosen by specifying the SUBSTUT operand on the COS macroinstruction for one COS entry, indicates the COS entry to be used when the COS specified in a session request is unknown. For more information about how the substitute COS works, see [“Substituting Class of Service Parameters” in “How session traffic is assigned to a specific route” on page 260](#).

Defining alias names

When defining resources in interconnected networks, you must reconcile any name conflicts among the connected networks. Following are the ways you can do this:

- Adopt a naming convention in all networks that ensures that all names are unique.
- Define alias name translation tables.

- Use the alias function in the session management exit routine.
- Implement network-qualified names.

The method that is best for you depends on how many names must be changed and how much control you have over the various networks.

To use an alias name translation facility, define all the logical units in the interconnected networks that are not already defined and their owning SSCPs to the NetView alias name translation facility and assign alias names. Then, if VTAM does not have a complete definition of a destination logical unit for a session setup request, it invokes the NetView alias name translation facility with the name of the logical unit. The NetView alias name translation facility returns the name of the SSCP defined as the owner of the logical unit, and the destination network ID and the logical unit real name in that destination network.

If the following conditions are true, the request for the session is sent to the owning SSCP:

- No predefined CDRSCs are used.
- VTAM has a session with the owning SSCP.
- Dynamic definition of cross-domain resources is allowed.

VTAM automatically creates a CDRSC definition for the logical unit.

If you have not installed a name translation facility, or if the name translation facility does not return the name of the owning SSCP, SSCP selection using the adjacent SSCP table is tried next.

Note: VTAM uses real names and alias names when referencing an LU. An alias name is the name of the LU as known in the network of its session partner. Because one LU can be in session with partners in more than one network, an LU can have multiple alias names. VTAM can obtain real and alias names through an alias translation application program (for example, NetView). If such an application program or the alias function of the session management exit routine is not in use, VTAM uses network qualifiers to define the real and alias names for logical units, in which case, all resource names must be unique. Therefore, no two LUs can have the same name, whether they are in the same network. However, if you are using network-qualified names, you can have duplicate resource names in different networks.

In general, when VTAM receives a session setup request and the alias name translation facility is active, VTAM invokes the alias name translation facility to translate LU names, Class of Service names, and the logon mode name. You can use the TRANSLAT start option to limit the scope of the translation functions for which VTAM invokes the NetView program.

The default for the CDRDYN start option is YES, enabling dynamic CDRSC definition. The CDRDYN start option overrides the CDRDYN operand on the host CDRM definition statement. CDRDYN must be YES for the VTAM in which the NetView alias name translation facility or the alias function of the session management exit routine is used.

If an alias name is defined to VTAM with a CDRSC definition (either predefined or dynamically defined), the DISPLAY operator command can be used to display the real name associated with that alias name (if a session exists with the resource and if VTAM has determined the real name).

You should have a convention for creating alias names. This alias naming convention should identify a name as an alias name. An operator often needs to know whether a resource name is an alias name or a real name. For example, the NetView program requires that the operator enter the real name of a resource when a display is requested.

Note: The generic name used by the alias function should not also be used by USERVAR.

As an example of name translation, [Figure 138 on page 469](#) shows a gateway NCP connecting two networks, NETA and NETB. Both networks contain an application program named IMS. A terminal user, TERMU, in NETA wants to log on to the IMS that is in NETB.

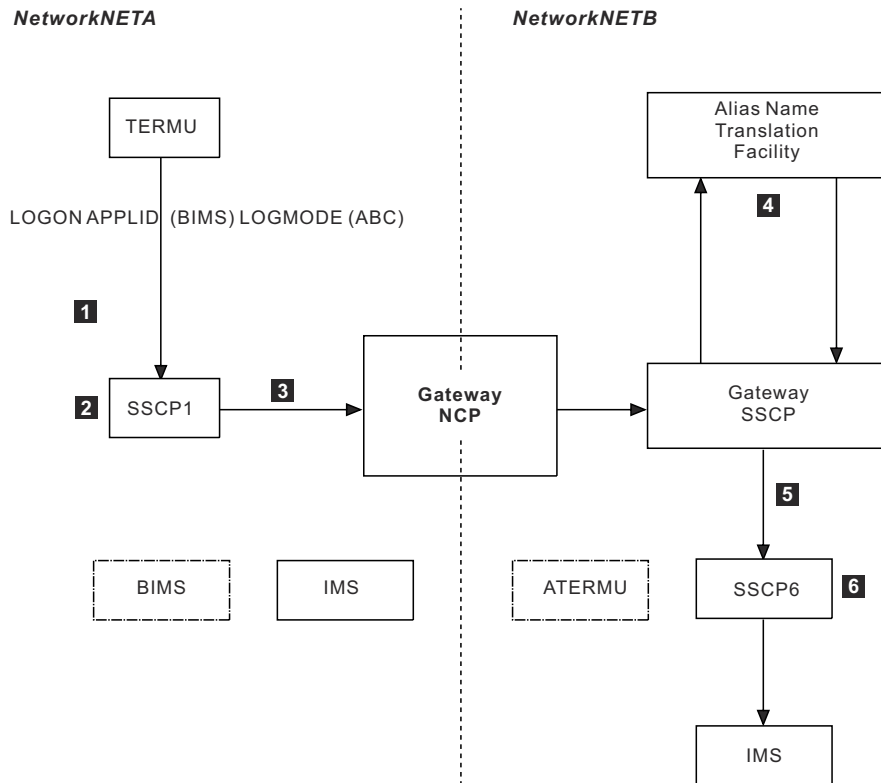


Figure 138. Example of name translation

The following steps show how name translation is used in setting up a session between TERMU in NETA and IMS in NETB. The numbers shown in Figure 138 on page 469 correspond to the following steps.

1. TERMU enters a LOGON command specifying BIMS, which is an alias name for IMS in NETB.
2. The LOGON command is processed by SSCP1, the SSCP that owns TERMU. The COS name FAST is obtained from the logon mode entry named ABC in the logon mode table associated with TERMU.
3. SSCP1 has a dynamic or predefined CDRSC definition without NETID for the resource named BIMS, which indicates that BIMS is owned by the gateway SSCP. SSCP1 therefore passes the request to the gateway VTAM.
4. Because the NetView alias name translation facility is active or because an alias function in the session management exit routine is used, the gateway VTAM supplies the following names for translation:
 - Origin network ID (NETA)
 - Origin LU (TERMU)
 - Alias name of destination LU (BIMS)
 - COS name (FAST)
 - Logon mode name (ABC)

What NetView or the session management alias function returns to the gateway VTAM depends on what user defined name translations are coded in the NetView program or are available from the session management alias function. In this example, the following names are returned:

- Destination network ID (NETB)
- Real name of destination LU (IMS)
- Name of owning SSCP for IMS (SSCP6)
- Alias name for TERMU (ATERMU)
- Equivalent COS name (QUICK)
- Equivalent logon mode name (XYZ)

5. A dynamic CDRSC definition is created in SSCP6 for TERMU using the alias name and is associated with the real name.

After the session is set up, it appears to SSCP1 that TERMU is in session with a CDRSC in NETA named BIMS. To SSCP6, it appears that IMS is in session with a CDRSC in NETB called ATERMU. However, if SSCP1 or SSCP6 were a VTAM with the NETID start option specified, SSCP1 or SSCP6 would also be aware that this is a cross-network session.

The rules for name translation are as follows:

- Logical unit names and logon mode names are translated only once. If class of service is provided on the session-setup request, it is translated once from the name known in the origin network to the name known in the destination network. As the response to the session setup request flows back to the origin host, it is translated in each network to determine the COS name as it is known in the adjacent network.
- If the destination LU name must be translated, at least one of the gateway VTAMs that controls the first gateway NCP in the session setup path must have an alias name translation facility.

Translation of the destination LU name is attempted in every gateway or gateway-capable SSCP that has an active name translation facility until the names are translated or until a gateway VTAM requests alias addresses from a gateway NCP. A gateway VTAM requests alias addresses if it has designated control of the gateway NCP or if it is the last gateway VTAM on the first gateway NCP. At that point the destination LU name is assumed to be real if it has not been translated.

- After the destination logical unit real name and network identifier are determined, VTAM still tries to translate the origin LU name at every gateway or gateway-capable SSCP along the session setup path that has an active name translation facility until the origin LU alias name is determined. Also, if VTAM knows that the destination LU name is real, it still invokes the alias name translation facility to try to translate the COS and logon mode names and to determine the owning SSCP.
- If you specify the owning SSCP in the alias tables, use the name specified on the SSCPNAME start option used by the VTAM that owns the resource.
- The logon mode name, Class of Service name, and the origin LU name can be translated at the gateway VTAM that translates the destination LU name or any subsequent gateway or gateway-capable SSCP along the session setup path.

Before VTAM Version 3 Release 2, the Class of Servicename for the destination network had to be translated at the same gateway VTAM that translates the destination LU name. If the destination LU were a real name, determined by CDRSC definition, the COS name had to be translated at the same host.

- NetView domain IDs (NCCFID) must be unique and cannot be assigned alias names.
- Alias name translation facility support for LU 6.2 is restricted. This facility can be used to determine the correct network identifier of the destination logical unit and the name of the owning SSCP. However, it cannot be used to translate logical unit names. Most peripheral nodes that use LU 6.2 communication protocol do not support alias name translation.
- The VTAM constants module limits the information to be translated by the alias name translation facility.

Resolving alias name requirements

An alias name is defined in a host to represent a logical unit, logon mode table, or Class of Service (COS) name in another network.

The following method can be used to determine which alias names are required and to distribute the required alias name definitions to the correct gateway SSCPs. The alias names you choose cannot match the names of any resources in your network.

1. Create a table listing all of your LUs, even if they do not have cross-network sessions. You need not list every resource in your network if it is obvious from your naming conventions that certain names cannot be duplicated in any network to which you are connecting.
2. List all of the resources in other networks with which you communicate. If you use an adjacent SSCP list and do not specify the destination network, include all resources in all networks in the second

column, even those with which you have no sessions. (This includes resources such as lines and physical units.) Otherwise, a resource might receive a session request that is intended for an LU that has the same name but is in a different network.

Table 47 on page 471 illustrates this process:

<i>Table 47. Network resource list example</i>		
All resources in your network (NETA)	All other network resources that communicate with owning network	Owning NETID for resources in other networks
ABLE	BAKER	NETY
CHARLIE	BAKER	NETZ
DOG	CHARLIE	NETX
FOX	EASY	NETX
EASY	JOAN	NETY
MARY	PETER	NETZ
ROGER	.	.

- Find any duplicates in the table you have created.

In the example shown in Table 47 on page 471, the duplicate names are CHARLIE and EASY. NETX.CHARLIE and NETX.EASY need alias names so you can distinguish them from resources in your network. Either NETY.BAKER or NETZ.BAKER needs an alias name so that you can distinguish them from each other.

For this example, the alias name definitions that you need to code for the NetView alias name translation facility are the following definition:

```
ORIGNET    NETA
LU         CHARLIE, NETX, ALCHARLI
LU         EASY, NETX, ALEASY
LU         BAKER, NETY, ALBAKER
```

ORIGNET always represents the network that knows the resource by its alias name. The network specified on the LU definition statement is always the network that knows the resource by its real name.

With these alias names replacing the real names, you now know every resource you communicate with by a unique name.

After every network goes through the preceding process, consolidate the results so that every network that is to request name translations has a complete set of alias name definitions for LUs. For example, if NETX has a NetView alias name translation facility, you should supply NETX with the following definition:

```
ORIGNET    NETA
LU         CHARLIE, NETX, ALCHARLI
LU         EASY, NETX, ALEASY
```

Similarly, if NETY has a NetView alias name translation facility, you should supply NETY with the following definition:

```
ORIGNET    NETA
LU         BAKER, NETY, ALBAKER
```

If your network has a NetView alias name translation facility, you should keep a copy of all the definitions you created. In addition, other networks might give you alias name definitions that they created to see your resources. For example, NETX might give you the following definition:

```
ORIGNET    NETX
LU         MARY, NETA, MARYX
```

```
LU EASY,NETA,EASYX
LU CHARLIE,NETA,CHARLIEX
```

By distributing the alias name definitions as indicated, you might find that several networks have the same definition. However, you are assured that LU names get translated at the first NetView alias name translation facility called, regardless of the direction of session setup.

Establishing and controlling SNA sessions

This section contains information to be considered for nonnative network type 2.1 connections. Also, to control cross-network sessions you can use the automatic logon function of VTAM and take advantage of certain NetView Performance Monitor functions.

Nonnative network type 2.1 connections

For nonnative network type 2.1 connections, consider the following when defining and specifying destination logical units.

Considerations for network nodes

In [Figure 139 on page 473](#), two AS/400 network nodes are attached to a VTAM network node by nonnative network type 2.1 connections. If two resources (one at each of the AS/400 network nodes) are to communicate with each other through the intermediate VTAM network node, at least one of the connections must be a LEN connection or APPN multiple network connectivity support must be implemented (see [“APPN multiple network connectivity” on page 75](#)).

Considerations for end nodes

In [Figure 139 on page 473](#), some resources are attached to the end node by nonnative network type 2.1 connections. If one of the resources initiates a session, the network node server caches the network-qualified name of the originating LU, and the resource can be found subsequently by the APPN network. However, APPN might not be able to locate these resources when attempting to locate them as destination LUs, unless the resources are predefined at the end node (with CDRSC definitions for all possible network IDs) and registered with the network node server. If all possible network-qualified names of the resource are known at the network node server, APPN searches can locate the destination LU. The use of a predefined alias name for these resources at the end node is not recommended.




Figure 139. Nonnative network type 2.1 connections

Specifying and defining the destination LU

When requesting a session with resources that are attached by nonnative network type 2.1 connections, the originating LU should specify the NETID of the adjacent node or no NETID. If the originating LU specifies the NETID of the adjacent node, the destination LU must be located in the same network as the adjacent node. If a session request does not include a NETID for the destination LU and a predefined alias CDRSC is defined for the destination LU, the resource can be located in any network.

Automatic logon

Using the LOGAPPL operand, you can define a cross-network controlling primary logical unit on any peripheral node (local or remote, switched or nonswitched). The effect is essentially the same as if a CDRSC within the same network were the controlling primary logical unit, except that the resulting session is cross-network rather than within the same network. The network-qualified name of the controlling primary logical unit can be specified with the LOGAPPL operand.

In a single-network environment, when the SSCP-SSCP session necessary for the establishment of a controlling primary logical unit session is not immediately available, VTAM stores information necessary for session setup. When the necessary SSCP-SSCP session becomes active, VTAM establishes the LU-LU session.

However, this process does not occur for cross-network sessions. You can establish cross-network controlling primary logical unit sessions through any number of intermediate SSCPs only if all of the required SSCP-SSCP sessions are active when the LU-LU session setup is attempted. In the case where the session setup path consists of more than one SSCP-SSCP session, the session setup proceeds as far as establishing the session control blocks at each active SSCP along the setup path. At some point along the path, an SSCP can determine that the session request cannot be routed any further because it has no active sessions with any SSCPs in its adjacent SSCP list for the specific destination logical unit. If this occurs, the request is rejected, the session control blocks are freed, and the operator is notified of the failure. After the necessary SSCP-SSCP sessions are active, the operator must try to reestablish the session. If all required SSCP-SSCP sessions exist, but the controlling application program is unavailable,

VTAM fails the initiation request until the application program is available. When the application program becomes available, VTAM automatically attempts to reestablish the session.

Operating VTAM

This section contains information about operation considerations in a multiple-network environment. If you are migrating to a multiple-network, take this information into consideration when making decisions about your operating procedures.

Using the NetView program for network management

The NetView program provides facilities for problem determination and network management in a gateway configuration. These facilities include:

- Information about the alias addresses, alias and real LU names, and owning SSCPs for LUs involved in cross-network sessions
- Session awareness data, such as session protocols
- Inquiry, on a session basis, as to whether a complete physical path exists between the controlling subarea nodes of the two session partners
- Display of the subarea nodes and transmission groups

The availability of these functions for a particular cross-network session depends on the NetView program being installed in the appropriate hosts. It also depends on the appropriate level of VTAM or NCP being installed in each subarea.

It is recommended that the NetView program be installed at every gateway VTAM host and at every host that controls resources that might have cross-network sessions. This ensures that the NetView program is able to provide the maximum amount of problem determination data for cross-network sessions.

Note: When Netview is active, some messages will be sent to the Netview instead of the MVS system console. See [z/OS Communications Server: SNA Messages](#) for further information about solicited and unsolicited messages and their destinations.

Application programs

For a multiple-network environment, you need to define the NetView program or another name translation facility and control application program security across networks.

Note: An application program can become a shadow resource if a CDRSC with the same name already exists when the major node containing the application program definition is activated. For more information about shadow resources, see [“Shadow resources”](#) on page 438.

Defining the NetView program

If you are running the NetView program, ensure that at least one application program major node contains definitions for the NetView program and the NetView alias name translation facility if you are using it. If you are using network-qualified names, it is not necessary to use the NetView alias name translation facility. Network-qualified names provide an alternative to using the facility for resolving duplicate names. The NetView alias name translation facility can be used in conjunction with network-qualified names.

If you want to use the alias name translation facility provided by the NetView program, by the NCCF licensed program, or by a similar user-written application program, include an APPL definition statement for that application program in the application program major node. The name on the APPL definition statement must be ALIASAPL, as shown in the following example:

```
ALIASAPL APPL AUTH=(CNM),PRTCT=password
```


Note: You can use a different name than ALIASAPL, but if you do, change and reassemble the CNM routing table. If you want to do this, see [z/OS Communications Server: SNA Customization](#) for more information.

The AUTH=CNM operand on the APPL definition statement means that this application program can use the communication network management (CNM) interface. This operand must be coded for all CNM application programs, including:

- The session monitor component of the NetView program
- The NetView program alias name translation facility
- The NLDM licensed program
- A user-written application program used for name translation

The following NetView program definition statements must also be coded:

- One ORIGNET definition statement for each network that requires alias name translation. The definition statements that follow this ORIGNET definition statement apply to the named network.
- One LU definition statement for each LU that is given an alias name.
- One COS definition statement for each Class of Service name that must be mapped to the name of an equivalent Class of Service used in another network.
- One MODE definition statement for each logon mode name that must be mapped to the name of an equivalent logon mode used in another network.

Chapter 19. Operating VTAM

Operating VTAM involves starting and stopping VTAM, activating and deactivating resources, and monitoring the domain. Although the operator usually performs these tasks, the system programmer must define the procedures that the operator uses when doing these tasks. These procedures are usually written by the system programmer in a locally-produced document called a *run book*.

A run book is a list of procedures for operators to follow. For example, you might want to list the start options that operators should use when starting different hosts. Following is an example of what might be in a run book:

```
After you activate S013H:
```

```
VARY NET,ACT,ID=S013H1  
VARY NET,ACT,ID=S013H2  
VARY NET,ACT,ID=S013H3  
VARY NET,ACT,ID=S013H4  
VARY NET,ACT,ID=S013H5  
VARY NET,ACT,ID=S013H6
```

Notice that this gives very specific information about your own network, such as the names of the nodes.

For operators who are learning about operating VTAM, see [z/OS Communications Server: SNA Operation](#).

The following list shows where to find information about operating VTAM.

- [“Starting the domain” on page 477](#)
- [“Activating resources” on page 481](#)
- [“Monitoring the domain” on page 486](#)
- [“Controlling the domain” on page 490](#)
- [“Deactivating resources” on page 491](#)
- [“Halting VTAM” on page 493](#)
- [“Canceling VTAM” on page 494](#)
- [“Automatic operations” on page 494](#)

Starting the domain

Use the START command to start VTAM. You can enter start options when you start VTAM, or you can specify a start list.

Procedure

1. Respond to message IEA101A during operating system IPL by entering RERP=xx, where xx represents the last two characters of the alternate ERP list name, if you want to use the resident error recovery procedures (ERP) option or the ERP option with an alternate ERP list.
2. Make available the necessary channel-attached devices using the MVS command VARY.
3. Use the VTAM start procedure.
4. Enter VTAM start options, if necessary.
5. Enter the logon manager START command, if necessary.

Configuration restart

You can use configuration restart to restart your network with the same status information before VTAM deactivation or failure. VTAM can maintain certain information about VTAM resources, such as status (active or inactive) and operator-specified values of activation and operating parameters. When restarting

VTAM after a halt or failure or when reactivating an individual major node after a deactivation or failure, VTAM can use this information to restore the resources to their status before the deactivation or failure.

Recording changes to the network configuration

When the host operating system is generated, VTAM definition statements are filed in the VTAM definition library. The first time a major node is activated, activation information is recorded in the resource definition table and in the NODELST file, if one exists. Any further activations or deactivations of major nodes are also recorded in the resource definition table and in the NODELST file (if any). Every time the VTAM operator changes the status of a minor node, the change is recorded in the resource definition table and in a configuration restart file, if one exists.

The network configuration can be restored to either its initial status or its status before failure or deactivation. The status being restored depends on whether you have defined configuration restart files and on how you use the CONFIG operand and the optional COLD or WARM operand when restarting the network. [Figure 140 on page 478](#) illustrates this process.

Using configuration restart warm start

As shown in [Figure 140 on page 478](#), if you specify the WARM operand when activating a major node, VTAM updates its resource definition table from information in the VSAM configuration restart file for the major node. VTAM then restores the minor nodes within the major node to their status before the failure or deactivation.

Also shown in [Figure 140 on page 478](#), if you specify the WARM start option when starting VTAM, VTAM can update its resource definition table from information in the NODELST file.

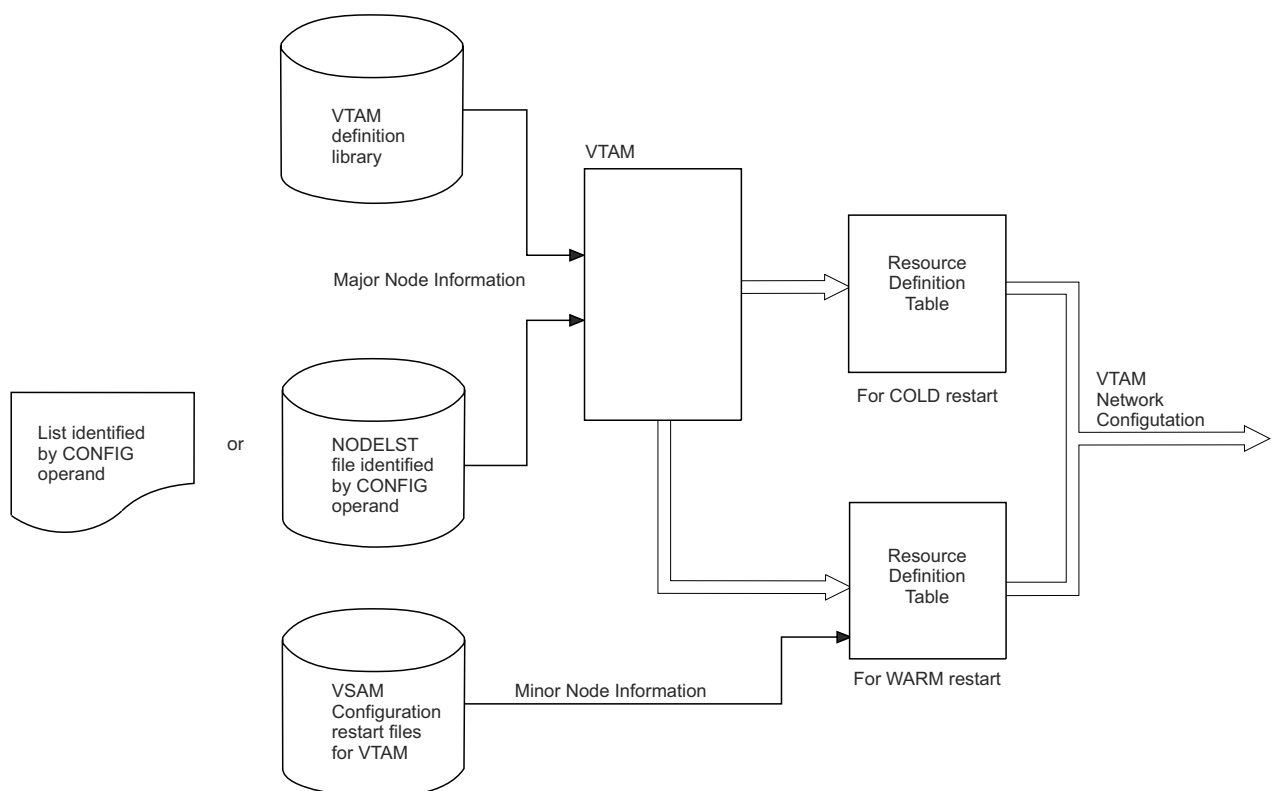


Figure 140. Restoring resource definitions with configuration restart and NODELST files

If you specified a NODELST file when the network was last started, the major nodes that were active and the DR files that were applied when the failure occurred can be reactivated by specifying the NODELST file name on the CONFIG operand during a warm start. If you used the VARY ACT,UPDATE technique of dynamic reconfiguration, major node definitions reflect dynamic reconfiguration changes, and the changes are reestablished when the major node is activated during NODELST processing. See [“Dynamic](#)

reconfiguration and change of operands” on page 174 for more information about the techniques that are available for dynamic reconfiguration.

Note: Dynamically defined independent logical units cannot be warm started.

For an example of how NODELST and configuration restart files are used, see “[Configuration restart example](#)” on page 479.

VTAM uses the VSAM configuration restart files to record changes to the named minor nodes within the major nodes (unnamed minor nodes are not checkpointed). If the VSAM files are defined, the VTAM operator can specify how VTAM is to use the files for configuration restart.

To associate VSAM configuration restart files with individual major nodes, code the CONFGDS and CONFGPW operands on the definition statements for the major nodes.

A NODELST file can be defined to record which major nodes and DR files are active in the network. A NODELST file is specified with the NODELST start option.

You can use either or both of these file types when using configuration restart.

If the user defines configuration restart files, VTAM provides two additional forms of recovery: manual switching to a backup host processor and manual switching to a backup communication controller.

If you do not associate a configuration restart file with a major node, or if the configuration restart file is empty, then the major node is not activated to its prior status.

If NODELST is in effect and you are doing a warm start of a configuration that has been dynamically reconfigured using the VARY DRDS technique, the DR statements that are processed again as part of the start process might fail because the nodes they specify became active earlier in the process.

Configuration restart example

Configuration restart files have been specified for the major nodes A, B, and C (CONFGA for major node A, CONFGB for major node B, CONFGC for major node C). Each major node is composed of four minor nodes (for example, major node A is composed of minor nodes A1, A2, A3, and A4). The following shows the minor nodes associated with each major node.

Major node A	Major node B	Major node C
A1 A2 A3 A4	B1 B2 B3 B4	C1 C2 C3 C4

Minor nodes B3 and B4 are to be initially active and the other minor nodes inactive.

Note: The initial status (active or inactive) can be specified for each minor node by coding the ISTATUS operand on the minor node definition statement.

The following chart shows the current contents of the configuration restart files for each major node.

CONFGA	CONFGB	CONFGC
empty	empty	empty

The VTAM operator enters the following start options on the START command:

```
CONFIG=XX,NODELST=ABC
```

The CONFIG start option indicates which major nodes are to be activated when VTAM is started. In this example, configuration list ATCCONXX contains major nodes A and C. The NODELST start option specifies that the VSAM file ABC is to be used to record which major nodes are active in the network.

Because of the above start options:

- VTAM activates major nodes A and C if it is able to open the NODELST file ABC.
- VTAM records the activation of major nodes A and C in file ABC.

Then to activate major node B, the operator enters:

```
VARY NET,ACT,ID=B
```

VTAM activates major node B and records its activation in ABC. The VSAM NODELST file ABC now contains the following active nodes:

- Major node A
- Major node B
- Major node C

VTAM also activates minor nodes B3 and B4 because ISTATUS=ACTIVE has been specified for them. Because B3 and B4 are defined as initially active, VTAM does not record the activation of B3 and B4 in CONFGB.

After the major nodes are active, the operator activates, deactivates, or modifies the status of minor nodes with VARY commands. Every time VTAM activates, deactivates, or modifies the status of a minor node, it records that action in the configuration restart file (if one exists) for the major node of which the minor node is a part.

```
VARY NET,ACT,ID=A2
VARY NET,ACT,ID=B3,LOGON=logapp1
VARY NET,INACT,ID=B3
VARY NET,ACT,ID=C4
```

As a result of the preceding commands, the status information stored in the configuration restart files can be represented as:

CONFGA	CONFGB	CONFGC
A2 active	B3 Status: LOGON B4 inactive	C4 active

If the network is now halted or if a failure occurs in the network, VTAM deactivates all the active major and minor nodes and closes the configuration restart files. The information needed to restore the network is in the VTAM definition statements and these files. When the VTAM operator restarts the network, start options can again be entered.

Start options specified	Major nodes restarted	Minor nodes restarted
CONFIG=XX,COLD	A and C	None
CONFIG=XX,WARM	A and C	A2 and C4
CONFIG=ABC,COLD	A,B, and C	B3 and B4
CONFIG=ABC,WARM	A,B, and C	A2, B3, and C4

Specifying CONFIG=XX,COLD or CONFIG=XX restores the network to its initial status while CONFIG=ABC,WARM restores the network to the status it had when VTAM stopped.

Note: COLD is assumed if neither COLD nor WARM is specified.

The operator can also specify another NODELST file:

```
CONFIG=ABC,WARM,NODELST=QRS
```

Or, the operator can continue using the one used before:

Information recorded by configuration restart

The following major nodes have the indicated information recorded:

Channel-attachment major node

- For each link, its status (active or inactive) and the channel unit address
- For each link station, its status (active or inactive)

Local SNA major node

- For each physical unit, its status (active or inactive) and its channel device address
- For each logical unit, its status (active or inactive) and the names specified in the LOGAPPL and LOGMODE operands and whether it is active or inactive
- For each logical unit, the cryptographic status (REQD, SEL, OPT, or NONE)

Switched SNA major node

- For each logical unit, its status (active or inactive) and the names specified in the LOGAPPL and LOGMODE operands
- For each physical unit, its status (active or inactive)
- For a physical unit with dial-out capability, the PATH=USE|NOUSE setting
- For each logical unit, the cryptographic status (REQD, SEL, OPT, or NONE)

Local non-SNA major node

Each channel-attached terminal, its status (active or inactive), its channel device address, and the names specified in the LOGAPPL and LOGMODE operands

CDRSC or CDRM major node

The status of each minor node (active or inactive)

Activating resources

You can initially activate resources by doing one of the following actions:

Specify the name of the resource in a configuration file

Do this to activate the resource when VTAM is initialized. For information about coding a configuration file, see [“Configuration lists” on page 28](#).

Use the VARY ACT command after VTAM is initialized

Enter a VARY ACT command with the ID operand specifying the name of the node to be activated.

To control activation of resources that are lower in the hierarchy, you can do one or both of the following actions:

Use the SCOPE operand on the VARY ACT command

If you code or use the default SCOPE=COMP on the VARY ACT command, all resources subordinate to this resource are activated if the ISTATUS operands for each are defaulted or coded as active.

SCOPE=ALL causes all resources to be activated regardless of the coding of the ISTATUS operand.

Code the ISTATUS operand on the resource definition statement

The ISTATUS operand controls the initial status of the resource. If you code or default to ISTATUS=ACTIVE, this resource is activated when its major node is activated.

If you make changes to your definition statements, deactivate and then reactivate the major nodes that are affected. Otherwise, the changes are not recorded.

Note: If you specify a warm start of VTAM, the status recorded in a configuration restart file overrides initial status. VTAM either activates a node or leaves a node inactive according to the status of the node when VTAM was last running. See [“Configuration restart example” on page 479](#) for additional information.

Order of activation

A resource cannot be activated unless all other resources above it in the hierarchy have been activated. For example, a logical unit cannot be activated unless the physical unit to which it is subordinate is already active. When a resource is deactivated, all of the resources subordinate to it in the hierarchy are also deactivated.

Although there is no logical hierarchy within the set of major nodes themselves or between resources in different major nodes, there is an obvious topological relationship among the major nodes representing subarea nodes in the network and the lines and link stations connecting them. Therefore, the activation of an NCP major node depends on the lines, link stations, and other subarea nodes that constitute the routes between VTAM and the NCP being active. Likewise, deactivation of an NCP major node or a line or link station within it can affect those other subarea nodes for which the resource being deactivated constitutes a part of the path from VTAM.

All of the elements of a communication path between two session endpoints must be active before a session can be established. In domains with more than one subarea, a path definition set must be activated before communication between subareas can be attempted. Major nodes are activated, and then their minor nodes are activated to enable sessions to be established for carrying user message traffic.

Although an activation command for a major node can be entered at any time after VTAM is initialized and before it is halted, the activation of NCP major nodes and external CDRM minor nodes requires fully operative explicit routes. If these routes are not available, the activations remain pending until the routes become available or until you deactivate the resources.

If automatic logons to a controlling application program have been defined for any logical units, all of the necessary application program major nodes must be activated before any of these logical units are activated (directly or indirectly). The logical units must be active before a session can begin.

The specific order in which you activate major nodes depends on how you have coded your network. For example, if you defined a terminal to automatically log on to an application program, activate the application program first. Following is a general guideline for activating your resources (the activation order may vary because of coding differences in your network):

1. Local resources, which includes the host CDRM, application programs, and other local resources. However, avoid activating any resources that will automatically log on to resources in other domains.
2. Routes, which include explicit and virtual route path definitions. Activate these paths before any connection to another subarea node (NCP or VTAM).
3. Peripheral resources, which include channels, NCPs, switched devices, cross-domain resource managers and cross-domain resources.

In multiple-domain networks, ADJSSCP tables should be activated before cross-domain resources, unless the dynamic adjacent SSCP function is specified.

When activating cross-domain resource managers, first activate the host cross-domain resource manager in each domain. This enables each host cross-domain resource manager to request sessions with and accept session requests from other known cross-domain resource managers.

After you have activated the host cross-domain resource manager, you can then activate any CDRM major nodes that represent other domains. VTAM then requests a session with the other SSCP. The request is rejected unless the other domain has activated both its host cross-domain resource manager and the external cross-domain resource manager representing the SSCP making the request.

Requests for sessions with cross-domain resources that are not defined are defined dynamically by the cross-domain resource manager, if you have authorized dynamic definition of cross-domain resources.

For autologon sessions, if a session request cannot be routed to its destination because an SSCP-SSCP session has not been established, the session fails and the operator receives a notice of the failure. After the needed SSCP-SSCP session is active, the session request is retried.

To cancel the autologon specification, a VARY NOLOGON command can be issued in the originating logical unit domain, or the logical unit can log on to a controlling LU. For nonautologon sessions, the operator receives notification of the failure, but the session attempt is not automatically retried.

Before logical units in one network can establish sessions with logical units in another network, successive sessions between gateway SSCPs connecting the networks must be established.

For an SSCP-SSCP session, if an NCP on the originating side of a gateway is not available, a session activation request remains pending. If an NCP is not available on the destination side, a session activation request fails.

In interconnected networks, the operator can test routes starting in a gateway NCP by using the NETID operand of the command with the ORIGIN operand.

Resources automatically activated by VTAM

Certain resources are automatically activated by VTAM. Some internally maintained resources are automatically activated when the message "VTAM INITIALIZATION COMPLETE" is issued. These resources can be displayed, but cannot be activated or deactivated by an operator. The following resources are automatically activated:

- VTAMSEG application program major node:
 - VTAM (or name from the CDRM definition statement for this VTAM)
 - ISTATA00
 - ISTNOP
 - ISTDCLU
 - ISTAPNCP
- VTAMSEG2 application program major node:
 - *?-* (model application program definition for Telnet server shared ACBs)

Note: The definition of the model application program for Telnet server shared ACB names cannot be displayed.

- ISTPUS PU (or name from HOSTPU start option) type 5 node
 - ISTGROUP
- ISTDILU predefined independent LU major node
- ISTADJCP adjacent CP major node
- ISTCDRDY dynamic cross-domain resource major node

Note: The ISTCDRDY major node can be deactivated and activated by an operator. For further information, see [“Dynamic definition of independent LUs” on page 191](#).

- ISTRTPMN rapid transport protocol major node
- ISTTRL transport resource list major node
- ISTLSXCF local SNA major node

Note: ISTLSXCF can also be deactivated and activated by the operator.

VTAM dynamically builds and activates transport resource list elements (TRLEs) within the ISTTRL major node for some TCP/IP communication interfaces. All of these TRLEs are created when needed, but cannot be deleted. These dynamic TRLEs are created with the following naming convention:

ISTTlsrs

TRLEs of this name are created when VTAM is started with either XCFINIT=YES (the default) or XCFINIT=DEFINE and another VTAM joins the XCF group (ISTXCF).

- *ls* is the two character &SYSCONE value of the VTAM on the local MVS image
- *rs* is the two character &SYSCONE value of the VTAM on the partner MVS image.

TCP/IP uses these TRLEs in one of the following situations:

- DYNAMICXCF is specified on the IPCONFIG or IPCONFIG6 statement and device or interface definitions are dynamically created to other VTAMs with XCF connectivity.
- DEVICE/LINK statements of type MPCPTP contain a device name that is the CPNAME or SSCPNAME of another VTAM with XCF connectivity.
- INTERFACE definitions of type MPCPTP6 contain a TRLENAM that is the CPNAME or SSCPNAME of another VTAM with XCF connectivity.

IUTOpfid

This TRLE is created when TCP/IP activates one of the following interfaces, which have the SMCD operand specified or taken as the default value, and Shared Memory Communications - Direct Memory Access (SMC-D) is enabled on the system:

- IPAQIDIO interface
- IPAQIDIO6 interface
- IPAQENET interface with CHPIDTYPE OSD
- IPAQENET6 interface with CHPIDTYPE OSD

The *pfid* value is discovered by VTAM during activation of the TRLE. No subchannels are associated with this TRLE.

IUTSAMEH

This TRLE is created for communication between multiple TCP/IP stacks on the same MVS image, and for communication between TCP/IP and VTAM for Enterprise Extender.

IUTIC6xx

This TRLE is created for TCP/IP when an IPv6 OSD interface with a defined physical network ID (PNetID) is activated and the TCP/IP stack has the AUTOIQDC GLOBALCONFIG option specified. The IQD CHPID associated with this TRLE is defined in the hardware configuration definition (HCD) as external bridge IQD function type and has a PNetID that matches the OSD CHPID. The value *xx* is the IQD CHPID number that is associated with this IQDC TRLE. Up to 10 subchannel addresses are allocated: one READ and one WRITE device, and eight DATAPATH devices.

IUTIQCxx

This TRLE is created for TCP/IP when an IPv4 OSD interface with a defined physical network ID (PNetID) is activated and the TCP/IP stack has the AUTOIQDC GLOBALCONFIG option specified. The IQD CHPID associated with this TRLE is defined in the hardware configuration definition (HCD) as external bridge IQD function type and has a PNetID that matches the OSD CHPID. The value *xx* is the IQD CHPID number that is associated with this IQDC TRLE. Up to 10 subchannel addresses are allocated: one READ and one WRITE device, and eight DATAPATH devices.

IUTIQDIO

This TRLE is created for TCP/IP dynamic XCF communications over HiperSockets devices. Up to 10 subchannel addresses are allocated: one READ and one WRITE device, and eight DATAPATH devices.

IUTIQDxx

This TRLE is created when TCP/IP activates a HiperSockets interface (defined by using either the DEVICE/LINK statements for IPAQIDIO or the IPv6 INTERFACE statement for IPAQIDIO6) with a CHPID parameter of *xx*. Up to 10 subchannel addresses are allocated: one READ and one WRITE device, and eight DATAPATH devices.

IUTIQXxx

This TRLE is created for TCP/IP dynamic Internal Queued Direct I/O extensions (IQDX) function IPv4 communications over HiperSockets devices that are connected to the intraensemble data network

(IEDN). The value *xx* is the OSX CHPID number that is associated with this IQDX TRLE. Up to 10 subchannel addresses are allocated: one READ and one WRITE device, and eight DATAPATH devices.

IUTIQ4xx

This TRLE is created when TCP/IP activates a HiperSockets interface (defined by using the IPv4 INTERFACE statement for IPAQIDIO) with a CHPID parameter of *xx*. Up to 10 subchannel addresses are allocated: one READ and one WRITE device, and eight DATAPATH devices.

IUTIQ6xx

This TRLE is created for TCP/IP dynamic IQDX IPv6 communications over HiperSockets devices that are connected to the intraensemble data network (IEDN). The value *xx* is the OSX CHPID number that is associated with this IQDX TRLE. Up to 10 subchannel addresses are allocated: one READ and one WRITE device, and eight DATAPATH devices.

IUTXT0xx

This TRLE is created when TCP/IP activates an MPCIPA INTERFACE with CHPIDTYPE OSX and a CHPID parameter of *xx*. Up to 19 subchannel addresses are allocated: one READ and one WRITE device, and 17 DATAPATH devices.

IUTMT0xx

This TRLE is created when TCP/IP activates a dynamically defined OSM interface, where VTAM assigned CHPID *xx* for this communication. Up to 11 subchannel addresses are allocated: one READ and one WRITE device, and nine DATAPATH devices.

IUTZCX4x

This TRLE is dynamically created when TCP/IP activates the z/OS Container Extensions (zCX) IPv4 interface named EZAZCX. The value *x* in the TRLE name denotes an internally assigned TCP/IP stack instance (1 through 8). A TCP/IP stack can support up to 255 unique IPv4 and IPv6 Linux Docker Container instances.

IUTZCX6x

This TRLE is dynamically created when TCP/IP activates the z/OS Container Extensions (zCX) IPv6 interface named EZ6ZCX. The value *x* in the TRLE name denotes an internally assigned TCP/IP stack instance (1 through 8). A TCP/IP stack can support up to 255 unique IPv4 and IPv6 Linux Docker Container instances.

IUTnpfid

This TRLE is created when TCP/IP activates an IPAQENET or IPAQENET6 interface with CHPIDTYPE OSD with Shared Memory Communications - RDMA (SMC-R) specified or taken as the default, and SMC-R is enabled on the system.

- For an IBM 10 GbE RoCE Express TRLE, the *npfid* value is derived from the PORTNUM and PFID values on the SMCR parameter of the GLOBALCONFIG statement in the TCP/IP profile. For example, IUT20018 indicates that the PORTNUM value is 2 and the PFID value is 0018. If PORTNUM is not specified, the default value is 1.
- For an IBM RoCE Express2 or RoCE Express3 TRLE, the *npfid* value is derived from the PFID value on the SMCR parameter of the GLOBALCONFIG statement or the IPAQENET INTERFACE statement in the TCP/IP profile and the port number that VTAM learns during activation of the RoCE Express2 or RoCE Express3 feature. The learned port number is used instead of any PORTNUM value specified on the GLOBALCONFIG SMCR statement. For example, IUT20153 indicates that the PFID value is 153 and that the learned port number is 2.

No subchannels are associated with this TRLE.

IUTtdddd

This TRLE is created when TCP/IP activates a CDLC, Hyperchannel, or CTC device.

- *t* identifies the type of device that is dynamically created:
 - C - TCP/IP CDLC
 - H - TCP/IP Hyperchannel
 - X - TCP/IP CTC
- *dddd* identifies the read device address for this device.

Activating application programs

Before an application program can use VTAM services, the application program minor node that defines the application program must be activated. An application program minor node is typically activated at VTAM startup by placing the name of the application program major node that contains the application program minor node into the VTAM configuration list.

After an application program minor node is activated, it is up to the application program to open its ACB and begin using VTAM services. The order in which the minor node is activated and the application program is started is not critical, but the minor node must be active when the application program attempts to open its ACB for VTAM.

If the operator cancels an application program, the operating system forces the application program disconnection from VTAM. For the operator to be able to cancel an application program that is active to VTAM, the operator must first associate the application program defined VTAM name with the application program job name.

The logon manager, which runs as a VTAM application, can be terminated with the MODIFY STOP command. For information about using the MODIFY STOP command, see [z/OS Communications Server: SNA Operation](#).

Monitoring the domain

VTAM automatically provides you with messages that contain information about the VTAM domain. If more information is required, various DISPLAY commands are available to request status information about domain resources, and to verify changes resulting from previous operator requests.

Using the DISPLAY command

In general, DISPLAY commands can be used to obtain status information about active major nodes and their minor nodes (active or inactive). Some DISPLAY commands can be used to obtain information about resources that are not nodes. For example, the DISPLAY BFRUSE command provides information about VTAM use of buffers, and the DISPLAY COS command provides information about the network Class of Service tables.

Monitoring I/O problems

You can change how long some VTAM I/O operations and internal procedures remain pending before VTAM notifies you. To do this, you use the IOINT start option, the MODIFY,VTAMOPTS command, or the MODIFY,IOPD command to change the I/O problem determination timeout interval. For example, if you use the command MODIFY,IOPD,IOINT=120, VTAM notifies you of I/O operations or internal procedures that remain pending for 120 seconds or more. After you are notified of a pending I/O operation, you can take any necessary action.

If you do not want to be notified of pending operations, you can specify IOINT=0. This deactivates the I/O problem determination function.

For information about the IOINT start option, see the [z/OS Communications Server: SNA Resource Definition Reference](#). For information about the MODIFY VTAMOPTS and the MODIFY IOPD commands, see [z/OS Communications Server: SNA Operation](#). See [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#) for more information about the IOPD trace facility.

Suppressing messages

To help you control the amount of information you receive when monitoring the network, VTAM allows some messages to be suppressed. Messages are divided into the following categories based on suppression level (in order, from the lowest suppression level to the highest):

1. Informational
2. Warning

3. Normal
4. Serious
5. Insuppressible

The categories to be suppressed can be specified, either with the SUPP start option or with the MODIFY SUPP command. By specifying one category, VTAM suppresses all messages in that category and all messages in lower categories. For example, if the operator issues MODIFY SUPP=NORM or if you code SUPP=NORM in the start options, all messages within the normal category and those below it, warning and informational, are suppressed.

Note: You cannot suppress messages classified as "insuppressible." Examples of insuppressible messages are those that solicit information from the operator, and messages that are responses to DISPLAY commands.

The MODIFY SUPP command changes only the categories of messages that are suppressed, not the suppression level of messages. You can change the suppression level of a message by using USS facilities.

For information about the suppression level of each VTAM message, see [z/OS Communications Server: SNA Messages](#).

Message flooding prevention

VTAM messages that could potentially flood the operator console can be suppressed by a message flooding prevention facility. This facility recognizes and filters duplicates occurring within a given time span, thus preventing critical information from being buried in a mass of repetitious information.

You can choose which messages are candidates for message flooding prevention by copying, renaming, and modifying VTAM default table (ISTMSFLD) to create your own message-flooding prevention table. Suppression criteria (variable text similarity) and suppression time intervals may be specified for each message in your table, and messages selected for suppression can also be suppressed from the hardcopy log. Message-flooding prevention support also provides the ability to modify a message flooding table dynamically.

For complete information about how to define a message-flooding prevention table, see the [z/OS Communications Server: SNA Resource Definition Reference](#). For information about modifying (MODIFY TABLE) and displaying (DISPLAY TABLE) the message-flooding prevention table, see [z/OS Communications Server: SNA Operation](#).

Other methods of controlling messages

Some VTAM messages can also be controlled using the following VTAM start options:

- ASIRFMSG
- CNNRTMSG
- DSIRFMSG
- DSPLYDEF
- DSPLYMAX
- ESIRFMSG
- FSIRFMSG
- HPRITMSG
- HPRPSMSG
- IOMSGLIM
- LSIRFMSG
- PLUALMSG
- RSIRFMSG
- SIRFMSG

- SLUALMSG

For information about these start options, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Displaying and testing routes

There are two VTAM commands used for displaying and testing routes, DISPLAY ROUTE and DISPLAY APING.

DISPLAY ROUTE command

The DISPLAY ROUTE command provides information about what routes are available between the VTAM host node and a destination subarea node. In addition, the TEST operand of the DISPLAY ROUTE command allows the operator to test the explicit routes between the host node and a destination subarea node for the routes' ability to transfer data between the subareas. The routes to be displayed can be specified by explicit route number, virtual route number, or Class of Service name. If the TEST operand is specified and the routes to be displayed are specified by virtual route number, the explicit routes tested are those to which the specified virtual routes are currently mapped. If the TEST operand is specified and the routes to be displayed are specified by Class of Service, the explicit routes tested are those to which the virtual routes in the specified Class of Service are currently mapped.

You can also test explicit routes starting in an NCP. The ORIGIN operand on the DISPLAY ROUTE command specifies the NCP in which the route starts. Although the originating node must be NCP Version 3 or later, the subarea at the other end can be any of the products that can coexist with VTAM Version 3 or subsequent releases.

DISPLAY APING command

Using the DISPLAY APING command, the network operator can:

- Verify connectivity with any LU 6.2 resource in the network
- Verify that another VTAM 4.3, or later, node is operational
- Check the performance of the network using a particular logon mode
- Display routing information to the destination node, if a new session is established for the APING transaction

For further information, see [“APING support” on page 337](#).

Defining operator messages and commands

You can modify certain messages and operator commands. These commands and messages are defined in unformatted system services (USS) tables and are called USS messages and USS commands. Whenever VTAM receives a USS command, it uses either an IBM-supplied default USS table or a user-written USS table to process it. Similarly, whenever VTAM is to send a USS message, it uses one of these tables to determine the message level (the message associated with this release of VTAM or a previous release), message text, and other characteristics of the message. The IBM-supplied operation-level USS table is named ISTINCNO.

By creating a supplementary USS table, you can change the text or other characteristics of a message or the syntax or default values for a command. You can create supplementary tables using USS macro instructions to redefine the VTAM commands or messages that you want to change. To use your USS table, specify the name of the table on the USSTAB start option. If you specify a user-defined table for operator commands, the ISTINCNO table is not used. If you specify a user-defined table for operator messages, the ISTINCNO table is used in conjunction with the user-defined USS table. For more information about USS tables, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Note: Because program operators depend on VTAM messages, changes to the operator messages could disrupt the functioning of a program operator. If you change an operator message, notify the responsible system programmer.

You can specify an operation-level USS table for a program operator using the SSCPFM and USSTAB operands of the program operator APPL definition statement.

Multiple console support (MCS) in VTAM

An operator can enter VTAM commands and receive messages at the system master console and at designated secondary consoles. To enter VTAM commands at a secondary console, the console must be authorized to accept commands from command group 1 (system control group) and command group 2 (I/O control group).

A VTAM command response goes to the console where the command was entered. Depending on their contents, VTAM messages use routing codes 1, 2, 4, 6, 8, or 10.

Therefore, a secondary console should be assigned routing codes based on the kinds of VTAM messages that are to be routed to that console.

To determine the routing codes used by particular messages, see [z/OS Communications Server: SNA Messages](#).

The operator can change command authorization and message routing using the MVS command VARY. The following command, for example, designates a secondary console where any VTAM command can be entered and where any VTAM message that has not been suppressed can be received:

```
VARY 01F,CONSOLE,AUTH=(SYS,IO)
```

The channel device name of the console is 01F. It is authorized to enter commands from Command Group 1 (SYS) and Command Group 2 (IO). This console receives all messages that have not been suppressed.

Details about using the VARY command to designate a secondary console are contained in the MVS reference document [z/OS MVS System Commands](#).

VTAM usage of IBM Function Registry for z/OS

The z/OS Communications Server (VTAM) product area has been registered to the IBM Function Registry for z/OS. VTAM has registered a SNA function with this z/OS service to record general usage statistics related to the SNA application programming interface (API) layer. Basic high-water mark counters are maintained for metrics related to VTAM applications and their associated sessions.

The function information and SNA usage statistics can be viewed using the following methods.

Using the system console

Use the following MVS command to display the current enablement state of all registered functions on the system console:

```
DISPLAY FXE VN=*,PN=*,PID=*,IID=*,FN=*,STATE
```

Example: The following command displays the state of all the functions registered for the z/OS Communications Server product for IBM:

```
DISPLAY FXE VN=IBM,PN='z/OS Communications Server',PID=*,IID=*,FN=*,STATE
```

For more information about the Display FXE command, see [Displaying function enablement state information \(FXE\) in z/OS MVS System Commands](#).

Using the IBM FXEPRINT utility function

The IBM FXEPRINT utility function generates a report that contains information about all the registered products. A JCL sample for using the FXEPRINT utility is included in SYS1.SAMPLIB member FXEPRNTJ and looks as follows:

```
//FXEPRINT EXEC PGM=FXEPRINT,TIME=1440,REGION=0K
//SYSOUT DD SYSOUT=A,DCB=(LRECL=80)
//SYSIN DD DUMMY
```

A sample display of the FXEPRINT job output follows:

```
-----
Vendor Name:          IBM
Vendor Description:   International Business Machines Corporation
Vendor Slot Path:     VS(1)
-----
Product Name:         z/OS Communications Server
Product Release:      03.01.00
Product ID:           HVT6310
Instance ID:          VTAMCS
Product Description:  VTAM
Product Slot Path:    VS(1) PS(1,-)
Product Parent:       IBM
Product Attributes:
  Attribute Name:      Last Updated
  Attribute Value:     07/26/22 11:32:10
-----
Function Name:         SNA - General
Function Description:  General SNA Information
Function Slot Path:    VS(1) PS(1,-) FS(1,AUTHONLY)
Function Parent:       z/OS Communications Server
Function Used:         YES
Function Enabled:      YES
Function Attributes:
  Attribute Name:      Maximum number of RAPI only applications
  Attribute Value:     0
  Attribute Name:      Maximum number of APPC capable applications
  Attribute Value:     0
  Attribute Name:      Maximum number of TSO applications
  Attribute Value:     2
  Attribute Name:      Maximum number of TN3270 applications
  Attribute Value:     0
  Attribute Name:      Maximum number of RAPI sessions
  Attribute Value:     0
  Attribute Name:      Maximum number of APPC sessions
  Attribute Value:     0
  Attribute Name:      Maximum number of TSO sessions
  Attribute Value:     1
  Attribute Name:      Maximum number of TN3270 sessions
  Attribute Value:     0
-----
```

Notes:

- The session counts are the high-water mark of the number of SNA session endpoints in this host. It should be noted that a session might have two endpoints in the same host, if so both endpoints will be counted.
- RAPI only applications are non-TSO, non-TN3270, and non-APPC=YES applications.
- APPC capable applications are applications whose APPL definition says APPC=YES. APPC capable applications can have RAPI sessions.
- RAPI sessions are the total number of non-TSO, and non-TN3270 sessions using VTAM's RAPI interface.
- APPC sessions are the 6.2 sessions established using the VTAM/APPC interface.
- Only non-internal applications and LU-LU sessions get counted. This explains the discrepancy observed in the number of LU-LU sessions displayed in the FXEPRINT output and the **d net,sessions,list=all** command.

Controlling the domain

Part of operating VTAM is the task of starting and stopping sessions under special conditions, controlling incoming and outgoing calls, and replacing or reassociating tables that affect how VTAM operates.

Establishing and terminating sessions with operator commands

A session is usually established by an LU requesting session establishment. However, you can establish a particular session or set of sessions using the VARY LOGON command, or you can establish a single session between a user and an application program using the LOGON operand of the VARY ACT command. If the logical unit you are activating is defined with a LOGAPPL on its LU definition statement, you can establish a session without using the LOGON operand.

A session is usually terminated by one of the session partners requesting termination. However, you can terminate a particular session or set of sessions without deactivating either of the session partners using the VARY TERM command. You can also terminate a session using the VARY INACT command for one of the logical units participating in the session. This command with a TYPE=IMMED or TYPE=FORCE operand breaks any sessions with the logical unit. (Note that TYPE=IMMED has the same effect as TYPE=UNCOND.)

Note: For cross-network sessions, when any SSCP-SSCP session along a cross-network session path of an LU-LU session is disrupted, the only termination requests that can be handled are those that are forced, such as VARY INACT with the FORCE operand or VARY TERM with the FORCE operand.

Dynamic table replacement

Using the MODIFY TABLE operator command, you can refresh, replace, or reassociate VTAM tables without halting and restarting VTAM. The following tables are affected by the MODIFY TABLE operator command:

- Associated LU tables
- Class of Service (COS) tables
- Interpret tables
- Logon mode tables
- Message-flooding prevention tables
- Model name tables
- Session awareness data filter tables
- USS tables

Before you can use the MODIFY TABLE command to replace a USS or interpret table, the new table must be assembled using the VTAM macro libraries.

Deactivating resources

This section provides an overview to deactivating VTAM resources with a VARY INACT command or a VARY REL command.

Order of deactivation

VTAM deactivates resources in the opposite order from activation. When VTAM receives a command to deactivate a resource, it first deactivates any resources subordinate to the named resource, moving from the bottom of the hierarchy upward, and then deactivates the named resource. Deactivating resources can be disruptive to existing LU-LU sessions. However, depending on the type of connection, a resource can be deactivated without disrupting the LU session if you use TYPE=GIVEBACK on the VARY INACT or the VARY REL command.

Deactivation of certain VTAM resources requires more attention than does their activation. Be aware of the changing connectivity of the domain and the effects on the rest of the network of NCP major node and channel-attachment major node deactivation. Also consider the effects of cross-subarea link and link station deactivations on the operation of the rest of the network. In particular, you should be aware of the impact of deactivating any of the following resources:

- An NCP through which other NCPs are attached to the network
- A link or link station, whether channel or SDLC, that provides the only connection between a subarea and the rest of the network

VTAM deactivates a cross-subarea link to a host processor when the channel-attachment major node is deactivated. Deactivating this link (a channel) destroys any sessions using a route on that link, possibly including the SSCP session with another subarea node. Also, VTAM might be an intermediate routing node

for cross-domain sessions that traverse the link. For this reason, care should be taken not to deactivate a cross-subarea link that could be supporting session traffic.

Note: When deactivating any channel-attachment major node, or any of its subordinate links or link stations, you might be deactivating a link that carries intermediate routing node traffic. Deactivating the link will disrupt these sessions.

When you enter a VARY INACT command for an NCP, VTAM does not deactivate any cross-domain link or link station unless CDLINK=INACT is specified on the VARY INACT command.

You can deactivate and reactivate CDRMs without disrupting active sessions. Use the SAVESESS operand on the VARY INACT command to specify whether active sessions should be terminated when you deactivate the CDRM. This operand enables you to keep active sessions when you use the FORCE or IMMED option of the VARY INACT command. However, SAVESESS specifies that all pending or queued sessions using the specified CDRM are terminated.

The loss of an SSCP-SSCP session (for example, as the result of a link failure) is not the same as external cross-domain resource manager deactivation and does not terminate existing LU-LU sessions. However, new sessions with logical units controlled by the SSCP that were lost cannot be started until the SSCP-SSCP session is reestablished.

Automatic deactivation

A VARY INACT command for an NCP major node can result in the automatic deactivation of cross-subarea links and link stations in adjacent subarea nodes (other NCPs or VTAM itself). For example, if the operator enters a command to deactivate an NCP, VTAM automatically deactivates any links and link stations that were automatically activated in adjacent nodes, provided these adjacent links and link stations were not also directly or indirectly activated.

VTAM does not automatically deactivate a multipoint subarea link, even if all physical units on the line have been deactivated.

Normal deactivation

Normal deactivation is the default for the VARY INACT command. When normal deactivation is requested, resources are not actually deactivated by VTAM until LU-LU sessions associated with the resources have been terminated. Active and pending-active sessions are not affected by normal deactivation of a resource. All queued requests for sessions involving the resources to be deactivated are discarded and the initiators are notified. No new requests for sessions with these resources are accepted.

Immediate deactivation

Immediate deactivation is specified with the TYPE=IMMED operand of the VARY INACT command. When immediate deactivation is requested, LU-LU sessions associated with the resources are effectively broken. All queued requests for sessions involving the resources to be deactivated are discarded, and the initiators are notified. No new requests for sessions involving these resources are accepted. All input and output operations for active LU-LU sessions are immediately halted, with possible loss of data. If the deactivated logical unit is in session with an application program, the application program is notified of the deactivation. The application program must complete the session termination for the deactivation to be completed.

Immediate deactivation provides tight control over the domain; normal deactivation provides less stringent control, but allows for a more orderly deactivation. Normal and immediate deactivations are not completed until the proper SNA requests and responses are exchanged between the VTAM SSCP and the affected resources. If conditions in the domain prevent the sending of the proper requests or responses, normal and immediate deactivation might not be successful.

Forced deactivation

Forced deactivation is specified with the TYPE=FORCE operand of the VARY INACT command. When forced deactivation is requested, all of the actions described for immediate deactivation apply. An

application program notified of the termination of a session with a deactivated logical unit might have to complete the session termination (depending on which user exits have been coded for the program) for the deactivation to be completed. Forced deactivation can be used when immediate and normal deactivation are not successful.

Forced reactivation

Forced reactivation is specified with the TYPE=REACT operand of the VARY INACT command. When forced reactivation is requested, VTAM performs a forced deactivation and then attempts to reactivate the resource. However, for lines, PUs, and LUs, VTAM does wait for responses from the resource. Forced reactivation of an NCP reactivates the SSCP-PU session, but has no effect on lines and link stations in adjacent subarea nodes. (In this case, the forced reactivation command should be entered for the adjacent line or link station if one of these resources is the source of the problem.) Forced reactivation does not apply to all types of resources.

Halting VTAM

You can halt VTAM using the HALT or HALT QUICK commands. When VTAM is halted, all of its resources are deactivated, and all sessions involving those resources are terminated. Additionally, if the VTAM being halted is an intermediate routing node, any sessions passing through that VTAM are disrupted. Sessions running under the IBM Network Routing Facility program might be disrupted if the VTAM host is halted. This is true in any situation where the session ends are not in the host subarea. If the HALT or HALT QUICK commands do not run properly, the operator can cancel VTAM.

The CDLINK operand for the HALT command is the same as for the VARY INACT command, except that the cross-subarea links to which it applies are determined when the HALT command is entered and are not redefined as the domain contracts. That is, the CDLINK operand applies to all cross-subarea SDLC links between NCPs that are cross-domain at the time the HALT command is entered.

Except for channel links, all other cross-subarea links (those that are entirely within the domain at the time the HALT command is entered) remain active during and after the halt, regardless of how CDLINK is specified. That is, their final status does not depend on whether the links and their link stations were automatically, directly, or indirectly activated. This allows sessions using that link to remain active. It also allows nondisruptive, simultaneous NCP deactivations, regardless of how the communication controllers are connected within the domain. Cross-subarea channel links are always deactivated during a VTAM halt.

The HALT command notifies application programs within the VTAM subarea of the domain shutdown and waits for them to close their ACBs. Intermediate routing node traffic passing through the VTAM being halted will be disrupted without notification.

New sessions are not permitted and new ACBs cannot be opened for any resources of the host that have received a HALT command. Except for this restriction, application programs of the host that have received a HALT command can continue their current operations. Any pending-active sessions are terminated, but active sessions remain active. After all application programs have stopped using VTAM services (that is, after they have closed their ACBs), processing is the same as for the HALT QUICK command.

During a normal halt, you can monitor the shutdown of the domain by displaying the status of application programs and other logical units. The VARY TERM command can be used to terminate sessions, and, if necessary, to speed the processing of the HALT command.

The HALT QUICK command can also be used at this time, if necessary, to speed halt processing. The HALT QUICK command notifies application programs of the domain shutdown and then proceeds with the equivalent of the VARY INACT,TYPE=IMMED command for each major node. All active LU-LU sessions are disrupted. After VTAM has received a HALT QUICK command, the only VTAM commands that can be used are DISPLAY commands, VARY TERM commands (to terminate sessions), and VARY INACT,TYPE=FORCE commands (to force the deactivation of resources).

If during HALT QUICK processing, after all the major nodes are inactive, any application programs have not yet closed their ACBs, VTAM displays a message listing the names of these application programs. Because VTAM cannot shut down the domain until all application programs have been disconnected from VTAM, you might want to speed up the halt process by canceling the jobs of application programs with

open ACBs. The HALT QUICK command can prevent an application program from using VTAM services, but VTAM cannot force the program to disconnect itself from VTAM. When the operator cancels an application program, the host operating system disconnects it from VTAM.

Canceling VTAM

It might be necessary to cancel VTAM if VTAM cannot be halted in any other way. VTAM should be canceled only if the HALT and HALT QUICK commands do not work properly. Canceling VTAM causes the immediate abnormal termination of VTAM without an orderly shutdown of the domain.

The HALT CANCEL command cancels VTAM. This command depends only on the proper functioning of the abnormal termination facilities on the host operating system. The HALT CANCEL command causes an abnormal termination with a system completion code of X'0A9'. When VTAM is canceled, no further I/O operations occur. Application programs using VTAM are notified of a VTAM shutdown through their TPEND exit routines, if possible. Data being transmitted on sessions might be lost. If TPEND cannot be scheduled for an application program, the application program is also abnormally terminated, if possible. In some circumstances, you might have to cancel some application programs.

Note: Any multinode persistent capable application programs will go into recovery pending state if the application is currently enabled for persistence.

Automatic operations

You can use special application programs to assist the VTAM operator in controlling a VTAM domain. These application programs are either program operators or communication network management (CNM) application programs.

Program operators

A program operator is an application program that is authorized to issue certain VTAM commands and receive VTAM messages. A program operator can be used to:

- Enter VTAM commands (except for commands to start or halt VTAM) from a terminal in the domain
- Monitor and control elements in the domain at program execution speed
- Define specialized commands (for example, to display the status of the entire domain)
- Reformat responses to VTAM commands (such as to reformat the status display of the entire domain to fit on a 3270 display screen)

Secondary and primary program operators

A program operator can have varying degrees of network management capability, depending on the options specified in its APPL definition statement. If you code AUTH=SPO (secondary program operator) or PPO (primary program operator), the application program can issue VTAM operator commands (except START and HALT) using the SENDCMD and RVCVMD macro instructions. The application program can then monitor and control the VTAM domain.

An application program that you specify as a primary program operator (PPO) receives all unsolicited messages, that is, all informational and error messages that are not replies to operator commands. If no application program that is designated as PPO is active when these messages occur, they are directed to the system console.

An application program that you specify as a secondary program operator (SPO) can receive solicited messages, that is, messages that are in reply to VTAM operator commands issued by that program operator only. You can, however, limit the number of messages that can be queued for a program operator by using the POAQLIM operand on the APPL definition statement.

CNM application programs

A communication network management (CNM) application program can request and receive information that might be useful in managing an SNA network. The application program can communicate with certain physical units in the network to obtain status information. [z/OS Communications Server: SNA Programming](#) describes in detail the macro instructions used to create a program operator.

CNM routing table

VTAM refers to a communication network management (CNM) routing table to determine which CNM application program is to receive an unsolicited network services request unit (NSRU) requiring further processing.

An application program can put its own procedure-related identifier (PRID) in each request sent to VTAM. When a reply to the request is returned, VTAM uses the PRID to route the reply to the application program. However, unsolicited request units contain network information, but they do not have a PRID. Therefore, VTAM needs the CNM routing table to determine which application program is to receive the unsolicited request unit.

VTAM provides default routing information for IBM CNM licensed programs, such as the NetView program. If you authorize a user-written application program to use the CNM interface, you should write a supplementary CNM routing table to route unsolicited requests to that application program.

For information about the request units for which VTAM provides default routing, see [z/OS Communications Server: SNA Customization](#).

NetView program and other CNM application programs

The NetView program can be used as a CNM application program or as a program operator. The hardware monitor and session monitor, which are components of the NetView program, can also be used as CNM application programs. Because the NetView program performs the functions of a program operator, code the program operator and CNM interface as authorized (AUTH operand on the APPL definition statement).

AUTH=CNM must be coded for the following application programs:

- NetView session monitor
- NetView alias name translation facility
- Network Logical Data Manager (NLDM) licensed program

Defining a CNM application program

To define a CNM application program, code AUTH=CNM on the APPL definition statement. The application program is then authorized to use the CNM interface to communicate with physical units in the domain, using embedded network services request units. Such an application program can use the CNM interface to gather error or link test information.

Collecting session awareness (SAW) data

Session awareness (SAW) data includes all information about sessions that is known by VTAM. A SAW data filter specifies which application programs receive SAW data and which do not. It can also control from which logical units the CNM application programs receive data.

Using the default SAW data filter

The default SAW data filter table allows data for all sessions to be passed across the CNM interface. You use the default filter only to reduce the amount of data that is passed to the session monitor, and therefore increase the amount of data that is sent to the CNM application programs.

ISTMGC10 in VTAMLIB contains the default SAW data filter. You can use the default table, modify the default table or dynamically change it using the MODIFY TABLE operator command, or you can replace it with one of your own.

NetView customers can use the default filter with the session monitor filter. The default filter then supplements the NetView filter. The session monitor filter controls the actual processing of NetView session awareness data and is still needed.

Coding your own SAW data filter

If you want more control over the amount of data that is filtered, you can use your own SAW data filter table, or you can modify the default table or dynamically change it using the MODIFY TABLE operator command.

You define a SAW data filter to VTAM using the KEEPMEM, KCLASS, and MAPSESS macros. Use the KCLASS and MAPSESS macros to define the session monitor filter. NetView customers can copy their existing filter definitions, modify them slightly, and create equivalent VTAM filters. The only change necessary is to place a KEEPMEM START macro at the beginning of the filter definitions and add a KEEPMEM STOP macro and an END statement at the end of the existing definitions. VTAM ignores any NetView operands that do not apply to the VTAM filter.

Following is a simple example of a SAW data filter:

```
ISTMGC10  KEEPMEM  START
SAW       KCLASS   SAW=YES
NOSAW     KCLASS   SAW=NO
M1        MAPSESS  KCLASS=SAW,PRI=SSCP1,SEC=*
M2        MAPSESS  KCLASS=SAW,PRI=IMS,SEC=*,PRINET=NETA
M3        MAPSESS  KCLASS=NOSAW,PRI=*,SEC=T3?7*,PRINET=*,SECTNET=NETB
          KEEPMEM  STOP
          END
```

The KCLASS macro instruction specifies whether to pass session awareness data over the CNM interface (SAW=YES|NO). The labels given to the KCLASS macro instructions are specified on the MAPSESS macro instruction. MAPSESS specifies a PLU-SLU session pair and specifies the KCLASS definition that VTAM uses to determine whether to pass session awareness data over the CNM interface.

The PLU and SLU are derived from the PRI and PRINET operands, and the SEC and SECTNET operands respectively. Each has a maximum length of eight characters. The special characters ? and * used in the pattern have the following meanings:

?

Matches any single letter or digit

*

Matches any letters or digits and can be used only as the last character in the pattern

The order of macro instructions is significant. VTAM examines the MAPSESS definitions from top to bottom when searching for a match for the PLU-SLU pair. The macro instructions in the simple example of a session awareness data filter table cause VTAM to:

- Pass session awareness data to the NetView session monitor for a session with a PLU named SSCP1 or IMS
- Ignore session awareness data for a session with an SLU whose name matches the T3?7* pattern if the PLU is not named SSCP1 or IMS
- Pass session awareness data for a session with IMS as the PLU, even if the SLU name matches the T3?7* pattern, because the M2 definition is before the M3 definition

Eliminating session initialization failure data

You can request that certain session initiation failure data not be sent to application programs collecting SAW data. By coding a SAW sense filter, you enable the SAW data to be filtered by sense code. Use the SNSCODE definition statement to specify either two or four bytes of the sense data code to be filtered. The following is a sample SAW sense filter:

```
SNSTBL1  VBUILD   TYPE=SNSFILTR
SENSE0   SNSCODE  SENSE=0857
SENSE1   SNSCODE  SENSE=087D0001
SENSE2   SNSCODE  SENSE=087D0002
```

SENSE3	SNSCODE	SENSE=087D0003
SENSE4	SNSCODE	SENSE=087D0004
SENSE5	SNSCODE	SENSE=8013

Based on this example, VTAM will not send:

- All sense data beginning with X'0857'
- The following sense data:
 - X'087D0001'
 - X'087D0002'
 - X'087D0003'
 - X'087D0004'

All other sense data beginning with X'087D' will be sent.

- All sense data beginning with X'8013'

Store the SAW sense filter in the VTAM definition library. Use the VARY ACT command to activate the SAW sense filter or to replace the current table. Only one SAW sense filter can be active in each VTAM host at the same time.

Chapter 20. Tuning VTAM for your environment

This section describes the task of tuning VTAM. It contains the following topics:

- [“Introduction to tuning” on page 499](#)
- [“Tuning tools” on page 500](#)
- [“Estimating the number of active sessions” on page 528](#)
- [“Common storage areas” on page 528](#)
- [“DISPLAY STORUSE pools” on page 529](#)
- [“Buffer pools” on page 536](#)
- [“Maximizing coattailing” on page 546](#)
- [“Session-level pacing tuning considerations” on page 555](#)

Introduction to tuning

Tuning is the process of balancing the network load among resources to eliminate congestion in any one resource. You need to tune VTAM to give you optimal service while using the least amount of resources. The objectives of tuning are to use storage in the host processor and in the communication controller more efficiently and to lessen the load on the host processor. For example, to use storage more efficiently, you should avoid allocating too many VTAM buffers in the host and choose a more appropriate buffer size in the communication controller.

The elements of a network can be regarded as a series of related storage spaces. A host processor provides storage for VTAM and application programs. Communication controllers contain storage for NCPs and related programs. Cluster controllers and programmable peripheral nodes also contain storage space.

When VTAM supports telecommunication services for your system, VTAM must use storage to build control blocks to keep track of sessions, blocks of data, and other information. If your system needs exceed your storage capacity, you might experience degraded response times. Similarly, an NCP can exceed its storage capacity if it receives more data from other parts of the network than it can send out.

VTAM uses storage for the following resources:

- VTAM modules
- VTAM buffer pools
- Tables representing major and minor nodes
- Tables representing routing capabilities of the SSCP
- Tables representing sessions controlled by the SSCP
- Temporary workspace used for VTAM operator commands and session establishment or termination requests

The amount of storage required for tables and buffer pools depends on the number of major nodes being defined and the number of sessions required by application programs. By increasing VTAM efficiency, you can free storage for other uses.

Tuning tools

To tune your environment, you need to be able to collect information about your network. You need to know how to gather and analyze tuning statistics, determine the amount of coattailing (when more than one piece of data is being transferred in or out of the host with a single I/O operation) taking place in your network, analyze slowdown conditions, and monitor your common storage areas.

VTAM provides operator commands and other facilities to help you gather this information. See [“Tuning tools” on page 500](#) for a description of the tools that you can use to gather data to tune your network.

Estimating active sessions

When VTAM provides session services for applications and logical units, VTAM allocates storage to handle information about where data is to be sent. You can help system performance by estimating the number of sessions that VTAM will have with a particular application program or logical unit. This way, VTAM can set aside an appropriate amount of space for handling sessions. For further information, see [“Estimating the number of active sessions” on page 528](#).

Common storage areas

VTAM uses common service area (CSA) to maintain some buffer pools and control blocks. Common storage areas are also used for other system needs, so you need to monitor VTAM use of common storage areas (for information see [“Monitoring common storage areas” on page 500](#)) and you might want to control VTAM use. For more information about limiting VTAM use of common storage areas, see [“Common storage areas” on page 528](#).

Buffer pools

VTAM uses buffer pools to control the handling of data: VTAM control blocks, I/O buffers, and channel programs that control the transmission of data. Each buffer pool handles storage for a different VTAM service; therefore, the needs of your network will determine which buffer pools are used the most. After you determine which buffer pools are the most important to your system, you may want to change the sizes of some of the buffer pools.

When you set a buffer pool size, you are reserving storage for use only by that buffer pool. If you specify buffer pools that are larger than necessary, you are setting aside storage that is not being used. The storage is being wasted when it could be used in other areas.

When you set a buffer pool size that is too small, VTAM will reach the buffer pool limit and will then dynamically allocate more space in the buffer pool. When the current need is satisfied, VTAM will then dynamically deallocate space in the buffer pool. Specifying small buffer pools conserves storage, but can cause frequent CPU use for expansion and contraction. For information about types of buffer pools and buffer pool allocation, see [“Buffer pools” on page 536](#).

Coattailing

Coattailing occurs when more than one piece of data is being transferred in or out of the host with a single channel input/output (I/O) operation. For example, without coattailing, each message read into VTAM generates a READ channel program. A READ channel program uses CPU cycles. If you can read more than one message into VTAM and generate only one read channel program, then coattailing is taking place and you are saving CPU cycles. Coattailing can provide greater throughput between a host and an SNA controller because fewer interrupts and instructions are processed for each message. For further information, see [“Maximizing coattailing” on page 546](#).

Tuning tools

This section describes the tools that you can use to gather data to tune your network.

Monitoring common storage areas

If availability of common service area (CSA) is a problem, you should monitor the VTAM buffer pools for unused storage. You can do this with the DISPLAY BFRUSE operator command or with the VTAM TYPE=SMS,ID=VTAMBUF trace. The output from this trace goes to the generalized trace facility (GTF).

The following example DISPLAY BFRUSE operator command report can be used to determine how much CSA is not being used.

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = BUFFER POOL DATA 860
IST920I I000      BUFF SIZE  619      EXP INCREMENT   504
IST921I          TIMES EXP   0        EXP/CONT THRESH 500 / *NA*
```

IST922I		CURR TOTAL	4002	CURR AVAILABLE	4001	
IST923I		MAX TOTAL	4002	MAX USED	5	
IST989I		EXP LIMIT	696846	BUFFS REQUESTED	0	
IST924I						
IST920I	BS00	BUFF SIZE	260	EXP INCREMENT	70	
IST921I		TIMES EXP	0	EXP/CONT THRESH	29	/ *NA*
IST922I		CURR TOTAL	70	CURR AVAILABLE	70	
IST923I		MAX TOTAL	70	MAX USED	0	
IST924I						
IST920I	LP00	BUFF SIZE	2032	EXP INCREMENT	6	
IST921I		TIMES EXP	0	EXP/CONT THRESH	2	/ *NA*
IST922I		CURR TOTAL	12	CURR AVAILABLE	10	
IST923I		MAX TOTAL	12	MAX USED	4	
IST924I						
IST920I	XD00	BUFF SIZE	697	EXP INCREMENT	5	
IST921I		TIMES EXP	0	EXP/CONT THRESH	4	/ *NA*
IST922I		CURR TOTAL	5	CURR AVAILABLE	5	
IST923I		MAX TOTAL	5	MAX USED	0	
IST924I						
IST920I	LF00	BUFF SIZE	120	EXP INCREMENT	30	
IST921I		TIMES EXP	0	EXP/CONT THRESH	1	/ *NA*
IST922I		CURR TOTAL	30	CURR AVAILABLE	26	
IST923I		MAX TOTAL	30	MAX USED	4	
IST924I						
IST920I	CRPL	BUFF SIZE	144	EXP INCREMENT	75	
IST921I		TIMES EXP	0	EXP/CONT THRESH	29	/ *NA*
IST922I		CURR TOTAL	75	CURR AVAILABLE	73	
IST923I		MAX TOTAL	75	MAX USED	2	
IST924I						
IST920I	SF00	BUFF SIZE	112	EXP INCREMENT	32	
IST921I		TIMES EXP	0	EXP/CONT THRESH	1	/ *NA*
IST922I		CURR TOTAL	32	CURR AVAILABLE	30	
IST923I		MAX TOTAL	32	MAX USED	2	
IST924I						
IST920I	SP00	BUFF SIZE	176	EXP INCREMENT	21	
IST921I		TIMES EXP	0	EXP/CONT THRESH	1	/ *NA*
IST922I		CURR TOTAL	21	CURR AVAILABLE	21	
IST923I		MAX TOTAL	21	MAX USED	0	
IST924I						
IST920I	AP00	BUFF SIZE	56	EXP INCREMENT	56	
IST921I		TIMES EXP	0	EXP/CONT THRESH	3	/ *NA*
IST922I		CURR TOTAL	56	CURR AVAILABLE	56	
IST923I		MAX TOTAL	56	MAX USED	0	
IST924I						
IST920I	TI00	BUFF SIZE	800	EXP INCREMENT	60	
IST921I		TIMES EXP	0	EXP/CONT THRESH	60	/ *NA*
IST922I		CURR TOTAL	26004	CURR AVAILABLE	26004	
IST923I		MAX TOTAL	26004	MAX USED	5	
IST924I						
IST920I	T100	BUFF SIZE	1340	EXP INCREMENT	36	
IST921I		TIMES EXP	0	EXP/CONT THRESH	60	/ *NA*
IST922I		CURR TOTAL	500	CURR AVAILABLE	500	
IST923I		MAX TOTAL	500	MAX USED	0	
IST924I						
IST920I	T200	BUFF SIZE	2028	EXP INCREMENT	300	
IST921I		TIMES EXP	0	EXP/CONT THRESH	60	/ *NA*
IST922I		CURR TOTAL	9000	CURR AVAILABLE	9000	
IST923I		MAX TOTAL	9000	MAX USED	0	
IST924I						
IST920I	CRA4	BUFF SIZE	4080	EXP INCREMENT	10	
IST921I		TIMES EXP	0	EXP/CONT THRESH	20	/ *NA*
IST922I		CURR TOTAL	100	CURR AVAILABLE	98	
IST923I		MAX TOTAL	100	MAX USED	4	
IST924I						
IST920I	CRA8	BUFF SIZE	8176	EXP INCREMENT	6	
IST921I		TIMES EXP	0	EXP/CONT THRESH	2	/ *NA*
IST922I		CURR TOTAL	12	CURR AVAILABLE	11	
IST923I		MAX TOTAL	12	MAX USED	4	
IST924I						
IST449I	CSALIMIT = 464565K, CURRENT = 40695K, MAXIMUM = 40695K					
IST790I	MAXIMUM CSA USED = 40695K					
IST1667I	SYSTEM CSA LIMIT = 516184K					
IST1831I	84% OF SYSTEM CSA STORAGE REMAINING = 433943K					
IST449I	CSA24 LIMIT = NOLIMIT, CURRENT = 59K, MAXIMUM = 59K					
IST790I	MAXIMUM CSA24 USED = 60K					
IST595I	IRNLIMIT = NOLIMIT, CURRENT = 0K, MAXIMUM = 0K					
IST981I	VTAM PRIVATE: CURRENT = 1352K, MAXIMUM USED = 1352K					
IST924I						
IST2403I	64-BIT	STORAGE TYPE		CURRENT	MAXIMUM	LIMIT
IST2404I		HVCOMMON		7M	7M	NOLIMIT
IST2405I		TRACE HVCOMMON		4M	4M	2048M
IST2413I		PRIVATE		22M	22M	NOLIMIT

```

IST2412I      FIXED HVCOMMON      11M      11M      NOLIMIT
IST2414I      FIXED PRIVATE       22M      22M      NOLIMIT
IST2415I      TOTAL FIXED        33M      33M      **NA**

IST924I -----
IST1565I CSA MODULES = 1796K
IST1565I CSA24 MODULES = 40K
IST1565I PRIVATE MODULES = 7700K
IST314I END

```

The general formula for determining the amount of CSA not being used for a given buffer pool follows:

$$\text{NOTUSED} = (\text{MAXTOTAL} - \text{MAXUSED}) \times \text{BUFFSIZE}$$

Applying this formula to the I/O buffer pool shown in the sample DISPLAY BFRUSE command report, the amount of CSA not being used follows:

$$\text{IOBUF unused} = (110 - 1) \times 334 = 36,406 \text{ bytes of fixed CSA}$$

The unused CSA for the other pools shown follows:

```

BSBUF unused = 7344 bytes of fixed CSA
LPBUF unused = 121920 bytes of pageable CSA
XDBUF unused = 6810 bytes of fixed CSA
LFBUF unused = 14160 bytes of fixed CSA
CRPL unused = 32960 bytes of pageable CSA
SFBUF unused = 6944 bytes of fixed CSA
SPBUF unused = 7200 bytes of pageable CSA
APBUF unused = 3136 bytes of fixed CSA
TIBUF unused = 37920 bytes of fixed CSA
CRA4 unused = 40800 bytes of pageable CSA
CRA8 unused = 98112 bytes of pageable CSA

```

Therefore, the conclusion is as follows:

```

TOTAL CSA unused = 406,512 bytes
TOTAL fixed      = 112,720 bytes

```

As this example illustrates, this procedure can help you to identify where CSA is not being used so that you can adjust your usage.

Analyzing slowdown

If you encounter a slowdown condition in your NCP, there are not enough buffers in the NCP to keep up with the amount of data being transferred to VTAM. An NCP slowdown condition causes a message to be issued to the operator. However, you can also examine the SLODN tuning statistic to analyze the frequency that an NCP reaches a CWALL threshold condition, which can be encountered after the NCP has entered slowdown.

Slowdown in a channel-attached SNA cluster controller means that the controller (for example, the 3174) is out of buffers. The SLODN tuning statistic identifies when an SNA controller enters this condition.

Slowdown for a multipath channel connection indicates the number of times that read subchannel has gone into slowdown. Slowdown occurs when there is a shortage of IOBUFs when deblocking a multipath channel (MPC) read buffer.

Gathering tuning information with the performance monitor interface

Some statistics pertinent to VTAM internal performance activity and resource utilization are collected by VTAM application programs known as monitors. After a monitor obtains the data, it can report information to its user. Monitors are useful in providing:

- Tuning and resource information
- Early warning of problems in the VTAM system
- Information useful in debugging
- Information for monitoring current usage and trends

- Historical data for long-term analysis

VTAM offers the performance monitor interface to facilitate the requesting and receiving of this data by monitors. This method is independent of other data collection techniques provided by VTAM, and does not affect their use.

In this collection process, VTAM provides raw data to a monitor, which is then responsible for providing:

- The operator interface to the user
- The user report mechanism
- Data analysis functions
- The management of any statistical database

Performance data can be selected by category to keep the overhead of statistical gathering at a minimum. See [z/OS Communications Server: SNA Customization](#) for the types of information that can be monitored.

For information about how to implement the performance monitor interface, see [z/OS Communications Server: SNA Programming](#).

Gathering tuning statistics

By using the TNSTAT start option or the MODIFY TNSTAT command, you can collect data that will help you set the proper values on resource definition operands that control VTAM I/O operations in your system.

You can use VTAM tuning statistics to gather information about the following connections:

- Enhanced QDIO (EQDIO)
- SNA controller
- Channel-to-channel
- Multipath channel
- TCP
- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE)

Restriction:

- You cannot use VTAM tuning statistics to gather information about internal shared memory (ISM) devices.
- You cannot use VTAM tuning statistics to gather information about zCX (EZA6CX or EZ6ZCX) interfaces.

Tip: You can obtain some tuning statistics for ISM interfaces by using the Netstat DEvlinks/-d report. The Netstat DEvlinks/-d report also provides some basic device statistics for zCX interfaces. For more information, see [Netstat DEvlinks/-d report in z/OS Communications Server: IP System Administrator's Commands](#).

System management facility (SMF) is required to record tuning statistics. Tuning statistics can optionally be displayed at the system console using the CNSL operand, and statistics are always recorded in the appropriate tuning statistics file. This file is an SMF data set. The tuning statistics record is SMF record type 50. The format depends upon the resource for which the tuning I/O operation is collected. The tables in this section show the formats that can be present in a tuning statistics record.

TNSTAT need not be specified in the VTAM start list to later activate tuning statistics.

No tuning statistics are provided for LANs connected through XCA lines.

Tuning statistics can be activated or deactivated for all devices simultaneously (global tuning statistics), and tuning statistics can also be activated or deactivated based on a TRLE name (TRLE tuning statistics). When a TRLE is first activated, the tuning statistics state for that TRLE is set to the global tuning statistics state. For instance, if global tuning statistics are active, TRLE tuning statistics are active for that TRLE.

Tuning input/output (I/O) operations for EQDIO interfaces

With EQDIO interfaces, z/OS Communications Server allocates a single subchannel associated with a Network Express OSH PCHID. With this architecture, each EQDIO interface operates using EQDIO queues: control and data. Gathering tuning statistics for an EQDIO interface will provide you with insight into the activity of these queues. In general, if you request tuning statistics for a given EQDIO interface (TRLE), you should see output records reporting the following information:

- Global Statistics: Contains global information that describes the number and type of each EQDIO queue. This record also contains statistics related to PCI interrupts that are relevant to the Control Read and Datapath Read queues.
- Control Read Queue: Contains data specific to the Control Read input queue.
- Control Write Queue: Contains data specific to the Control Write output queue.
- Datapath Write Queues: Contains data related to write operations of a specific WLM priority between 1 and 4. Four sequential Datapath Write records may be produced.
- Datapath Read Queues: Contains data related to inbound read operations specific to one of the customized Inbound Workload Queue (IWQ) input queues. Up to seven sequential Datapath Read records may be produced.

Online tuning statistics report for EQDIO interfaces

This is an example of a tuning statistics report for a Network Express (EQDIO) interface. The message output for each Network Express interface begins with the Global statistics record followed by sections that represent each of the EQDIO input/output queues (See above for details). For a description of these fields, see message IST1230I in z/OS Communications Server: SNA Messages.

```
IST1230I TIME      = 14362989   DATE      = 24059      ID      = IUTE2B71 059
IST2488I DEV       = 2B71       QID/TYPE  = GLOBAL
IST1719I PCIREALO  =           0 PCIREAL   =          101
IST1720I PCIVIRTO  =           0 PCIVIRT   =           0
IST1750I PCITHRSO  =           0 PCITHRSH  =           0
IST1751I PCIUNPRO  =           0 PCIUNPRD  =           0
IST1752I RPROCDEO  =           0 RPROCDEF  =           0
IST1753I RREPLDEO  =           0 RREPLDEF  =           0
IST924I
-----
IST2488I DEV       = 2B71       QID/TYPE  = CTRL/0 (READ)
IST1754I NOREADSO  =           0 NOREADS   =           0
IST1722I PACKCNTO  =           0 PACKCNT   =           0
IST1236I BYTECNTO  =           0 BYTECNT   =           0
IST924I
-----
IST2488I DEV       = 2B71       QID/TYPE  = CTRL/1 (WRITE)
IST1722I PACKCNTO  =           0 PACKCNT   =           0
IST1236I BYTECNTO  =           0 BYTECNT   =           0
IST924I
-----
IST2488I DEV       = 2B71       QID/TYPE  = WR/2 PRIORITY1
IST2482I SBPEMAX   =           0 SBPEAVG   =           0
IST1756I QDPHMAX   =           0 QDPHAVG   =           0
IST2483I SETCNTO   =           0 SETCNT   =           0
IST2484I SETICNTO  =           0 SETICNT   =           0
IST2485I SETFCNTO  =           0 SETFCNT   =           0
IST2486I SBPECNTO  =           0 SBPECNT   =           0
IST1722I PACKCNTO  =           0 PACKCNT   =           0
IST1236I BYTECNTO  =           0 BYTECNT   =           0
IST1810I PKTIQDO   =           0 PKTIQD    =           0
IST1811I BYTIQDO   =           0 BYTIQD    =           0
IST924I
-----
IST2488I DEV       = 2B71       QID/TYPE  = WR/3 PRIORITY2
IST2482I SBPEMAX   =           0 SBPEAVG   =           0
IST1756I QDPHMAX   =           0 QDPHAVG   =           0
IST2483I SETCNTO   =           0 SETCNT   =           0
IST2484I SETICNTO  =           0 SETICNT   =           0
IST2485I SETFCNTO  =           0 SETFCNT   =           0
IST2486I SBPECNTO  =           0 SBPECNT   =           0
IST1722I PACKCNTO  =           0 PACKCNT   =           0
IST1236I BYTECNTO  =           0 BYTECNT   =           0
IST1810I PKTIQDO   =           0 PKTIQD    =           0
IST1811I BYTIQDO   =           0 BYTIQD    =           0
IST924I
-----
IST2488I DEV       = 2B71       QID/TYPE  = WR/4 PRIORITY3
IST2482I SBPEMAX   =           0 SBPEAVG   =           0
IST1756I QDPHMAX   =           0 QDPHAVG   =           0
```

```

IST2483I SETCNT =          0 SETCNT =          0
IST2484I SETICNT =          0 SETICNT =          0
IST2485I SETFCNT =          0 SETFCNT =          0
IST2486I SBPECNT =          0 SBPECNT =          0
IST1722I PACKCNT =          0 PACKCNT =          0
IST1236I BYTECNT =          0 BYTECNT =          0
IST1810I PKTIQD =          0 PKTIQD =          0
IST1811I BYTIQD =          0 BYTIQD =          0
IST924I -----
IST2488I DEV      = 2B71      QID/TYPE = WR/5 PRIORITY4
IST2482I SBPEMAX =          0 SBPEAVG =          0
IST1756I QDPHMAX =          0 QDPHAVG =          0
IST2483I SETCNT =          0 SETCNT =          101
IST2484I SETICNT =          0 SETICNT =          0
IST2485I SETFCNT =          0 SETFCNT =          0
IST2486I SBPECNT =          0 SBPECNT =          101
IST1722I PACKCNT =          0 PACKCNT =          101
IST1236I BYTECNT =          0 BYTECNT =          30502
IST1810I PKTIQD =          0 PKTIQD =          0
IST1811I BYTIQD =          0 BYTIQD =          0
IST924I -----
IST2488I DEV      = 2B71      QID/TYPE = RD/6 (PRIMARY)
IST2482I SBPEMAX =          0 SBPEAVG =          0
IST2487I LOWREAD =          0 LOWREAD =          0
IST2185I FRINVCT =          0 FRINVCT =          0
IST2486I SBPECNT =          0 SBPECNT =          0
IST1722I PACKCNT =          0 PACKCNT =          0
IST1236I BYTECNT =          0 BYTECNT =          0
IST1810I PKTIQD =          0 PKTIQD =          0
IST1811I BYTIQD =          0 BYTIQD =          0
IST924I -----
IST2488I DEV      = 2B71      QID/TYPE = RD/7 (BULKDATA)
IST2482I SBPEMAX =          0 SBPEAVG =          0
IST2487I LOWREAD =          0 LOWREAD =          0
IST2185I FRINVCT =          0 FRINVCT =          0
IST2486I SBPECNT =          0 SBPECNT =          0
IST1722I PACKCNT =          0 PACKCNT =          0
IST1236I BYTECNT =          0 BYTECNT =          0
IST1810I PKTIQD =          0 PKTIQD =          0
IST1811I BYTIQD =          0 BYTIQD =          0
IST314I END

```

Tuning statistics file report for EQDIO interfaces

Statistics that are produced are always written to the appropriate tuning statistics file in the format shown in the tables. For each EQDIO interface, a minimum of nine records are written:

- Global
- Control Read
- Control Write
- Write – Priority 1
- Write - Priority 2
- Write – Priority 3
- Write – Priority 4
- Read - Primary Input Queue
- Read – Bulk Input Queue

EQDIO Global statistics record

This record contains global statistics related to PCI interrupts that are relevant to the Control Read and Datapath Read queues.

<i>Table 48. EQDIO Global Statistics record</i>			
Offset	Length	Format	Description
0	2	Binary	Record length

Table 48. EQDIO Global Statistics record (continued)

Offset	Length	Format	Description
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	packed	Date: 01YYDDDDF Date: 00YYDDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	TRLE name
26	4	Binary	Reserved
30	8	Binary	MVS token
38-45	8	Binary	Reserved
46	4	Binary	Device address (DEV)
50	4	Binary	Reserved
54	4	EBCDIC	Queue Type = 'GLBL'
58	4	Binary	Reserved
62	1	Binary	Extension length (including this field)
63	1	Binary	EQDIO Attachment type = X'FF'
64	2	EBCDIC	EQDIO version = X'F0F6'
66	2	Binary	Reserved
68	4	Binary	Number of EQDIO control READ queues
72	4	Binary	Number of EQDIO control WRITE queues
76	4	Binary	Number of EQDIO datapath WRITE queues
80	4	Binary	Number of EQDIO datapath READ queues
84	4	Binary	PCI real interrupts overflow
88	4	Binary	PCI real interrupts count
92	4	Binary	PCI virtual interrupts overflow
96	4	Binary	PCI virtual interrupts count
100	4	Binary	PCI threshold overflow
104	4	Binary	PCI threshold count
108	4	Binary	PCI unproductive overflow
112	4	Binary	PCI unproductive count
116	4	Binary	Read processing deferrals overflow
120	4	Binary	Read processing deferrals count
124	4	Binary	Read replenishment deferrals overflow
128	4	Binary	Read replenishment deferrals count
132-277	146	Binary	Reserved

EQDIO Control Read record

This record contains data specific to the Control Read input queue

<i>Table 49. EQDIO Control Read record</i>			
Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	packed	Date: 01YYDDDF Date: 00YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	TRLE name
26	4	Binary	Reserved
30	8	Binary	MVS token
38-45	8	Binary	Reserved
46	4	Binary	Device address (DEV)
50	4	Binary	Reserved
54	4	EBCDIC	Queue Type = 'READ'
58	4	Binary	Reserved
62	1	Binary	Extension length (including this field)
63	1	Binary	EQDIO Attachment type = X'01'
64	2	EBCDIC	EQDIO Version = X'F0F6'
66	2	Binary	Reserved
68	4	Binary	Reads exhausted overflow
72	4	Binary	Reads exhausted count
76	4	Binary	Packet count overflow
80	4	Binary	Packet count
84	4	Binary	Byte count overflow
88	4	Binary	Byte count
92 - 277	186	Binary	Reserved

EQDIO Control Write record

This record contains data specific to the Control Read input queue

<i>Table 50. EQDIO Control Write record</i>			
Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type

Table 50. EQDIO Control Write record (continued)

Offset	Length	Format	Description
6	4	Binary	Time record moved to buffer
10	4	packed	Date: 01YYDDDF Date: 00YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	TRLE name
26	4	Binary	Reserved
30	8	Binary	MVS token
38-45	8	Binary	Reserved
46	4	Binary	Device address (DEV)
50	4	Binary	Reserved
54	4	EBCDIC	Queue Type = 'WRIT'
58	4	Binary	Reserved
62	1	Binary	Extension length (including this field)
63	1	Binary	EQDIO Attachment type = X'02'
64	2	EBCDIC	EQDIO Version = X'F0F6'
66	2	Binary	Reserved
68	4	Binary	Packet count overflow
72	4	Binary	Packet count
76	4	Binary	Byte count overflow
80	4	Binary	Byte count
84-277	194	Binary	Reserved

EQDIO Datapath Read record

This record contains data related to inbound read operations specific to one of the customized Inbound Workload Queue (IWQ) input queues.

Table 51. EQDIO Datapath Read record

Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	packed	Date: 01YYDDDF Date: 00YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	TRLE name
26	4	Binary	Reserved
30	8	Binary	MVS token

Table 51. EQDIO Datapath Read record (continued)

Offset	Length	Format	Description
38-45	8	Binary	Reserved
46	4	Binary	Device address (DEV)
50	4	Binary	Reserved
54	5	EBCDIC	Queue Type / QID = 'RD/xy'
59	3	Binary	Reserved
62	1	Binary	Extension length (including this field)
63	1	Binary	EQDIO Attachment type = X'04'
64	2	EBCDIC	EQDIO version = X'F0F6'
66	2	Binary	Reserved
68	4	Binary	Maximum SBPEs
72	4	Binary	Average SBPEs
76	4	Binary	Maximum EPCQs
80	4	Binary	Average EPCQs
84	8	EBCDIC	READ queue name
92	4	Binary	Reads exhausted overflow
96	4	Binary	Reads exhausted count
100	4	Binary	Packet count overflow
104	4	Binary	Packet count
108	4	Binary	Byte count overflow
112	4	Binary	Byte count
116	4	Binary	EPCQs exhausted overflow
120	4	Binary	EPCQs exhausted count
124	4	Binary	SBPE count overflow
128	4	Binary	SBPE count
132	4	Binary	EPCQ count overflow
136	4	Binary	EPCQ count
140	4	Binary	Frame invalidation packet count overflow
144	4	Binary	Frame invalidation packet count
148	4	Binary	Accelerated packet count overflow
152	4	Binary	Accelerated packet count
156	4	Binary	Accelerated byte count overflow
160	4	Binary	Accelerated byte count
164-277	114	Binary	Reserved

EQDIO Datapath Write record

This record contains data related to write operations of a specific WLM priority between 1 and 4.

<i>Table 52. EQDIO Datapath Write record</i>			
Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	packed	Date: 01YYDDDF Date: 00YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	TRLE name
26	4	Binary	Reserved
30	8	Binary	MVS token
38-45	8	Binary	Reserved
46	4	Binary	Device address (DEV)
50	4	Binary	Reserved
54	4	EBCDIC	Queue Type / Queue ID = 'WR/x'
58	1	Binary	Outbound write priority – 1 through 4
59	3	Binary	Reserved
62	1	Binary	Extension length (including this field)
63	1	Binary	EQDIO Attachment type = X'03'
64	2	EBCDIC	EQDIO version = X'F0F6'
66	2	Binary	Reserved
68	4	Binary	Maximum SBPEs
72	4	Binary	Average SBPEs
76	4	Binary	Maximum queue depth
80	4	Binary	Average queue depth
84	4	Binary	Packet count overflow
88	4	Binary	Packet count
92	4	Binary	Byte count overflow
96	4	Binary	Byte count
100	4	Binary	SETCP count overflow
104	4	Binary	SETCP count
108	4	Binary	SETCP init given overflow
112	4	Binary	SETCP init given count
116	4	Binary	SETCP force overflow

Table 52. EQDIO Datapath Write record (continued)

Offset	Length	Format	Description
120	4	Binary	SETCP force count
124	4	Binary	SBPE overflow
128	4	Binary	SBPE count
132	4	Binary	Accelerated packet count overflow
136	4	Binary	Accelerated packet count
140	4	Binary	Accelerated byte count overflow
144	4	Binary	Accelerated byte count
148-274	130	Binary	Reserved

Tuning input/output (I/O) operations for SNA controllers

VTAM performs I/O operations to channel-attached communication controllers and cluster controllers.

You can use VTAM tuning statistics to adjust operands in NCP and VTAM definition statements to make your system run more efficiently.

VTAM uses channel programs to send data to SNA controllers. The amount of data that VTAM can read in one operation depends on the number of buffers used by a read channel program and on the size of each buffer.

Each tuning statistics record contains information about the state of data-transfer operations between VTAM and one channel-attached SNA controller (communication or cluster controller). Each record contains statistics that cover the time period since the last tuning statistics record was written for that controller or channel-to-channel connection.

Online tuning statistics report for SNA controllers

Following is an example of a tuning statistics report for an SNA controller that can appear at the VTAM operator console:

```

IST440I    TIME = 14531106    DATE = 07093    ID = 0321-L
IST441I    DLRMAX = 3        CHWR = 109672    CHRD = 116797
IST442I    ATTN = 82624      RDATN = 73771    IPDU = 469283
IST443I    OPDU = 395495     RDBUF = 645489    SLODN = 0
IST1568I   INLP = 4202      ONLP = 4170     BFNLP = 853639
IST314I   END

```

For a description of these fields, see message IST440I in [z/OS Communications Server: SNA Messages](#).

Tuning statistics file report for SNA controllers

If statistics are produced, they are always written to the appropriate tuning statistics file in the format shown in Table 53 on page 511.

Table 53. Record format for SNA controller

Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer

Table 53. Record format for SNA controller (continued)

Offset	Length	Format	Description
10	4	Packed	Date: 01YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	Controller name
26	4	Binary	Dump load restart requests (DLRMAX)
30	4	Binary	Write channel program count (CHWR)
34	4	Binary	Read channel program count (CHRD)
38	4	Binary	Total attention interrupts received (ATTN)
42	4	Binary	Attentions on ends of READs (RDATN)
46	4	Binary	Number of PDUs inbound (IPIU)
50	4	Binary	Number of PDUs outbound (OPIU)
54	4	Binary	Total read buffers used (RDBUF)
58	4	Binary	Number of slowdowns (SLODN)
62	1	Binary	Extension length (including this field)
63	1	Binary	Reserved
64	2	EBCDIC	Version = X'F0F4'
66	4	Binary	Reserved
70	4	Binary	Reserved
74	4	Binary	Reserved
78	4	Binary	Reserved
82	4	Binary	Number of inbound NLPs (INLP)
86	4	Binary	Number of outbound NLPs (ONLP)
90	4	Binary	Number of bytes read from inbound NLPs (BFNLP)
94	4	Binary	Reserved
98	4	Binary	Reserved
102	4	Binary	Reserved
106	4	Binary	Reserved
110	4	Binary	Reserved

Tuning I/O operations for channel-to-channel connections

VTAM performs I/O operations across a channel-to-channel link. By analyzing statistics provided by VTAM, you can select I/O buffer sizes, data transfer delay, and other options that can improve performance of these I/O operations.

VTAM uses channel programs to send data to other hosts. The amount of data that VTAM can read in one operation depends on the number of buffers used by a read channel program and on the size of each buffer.

Each tuning statistics record contains information about the state of data-transfer operations between two VTAMs (using channel-to-channel adapters). Each record contains statistics that cover the time period since the last tuning statistics record was written for that controller or channel-to-channel connection.

Online tuning statistics for channel-to-channel adapters

Following is an example of a tuning statistics report for a channel-to-channel attachment:

```
IST577I    TIME  = 10214650    DATE = 94010    ID = LINEBC0
IST578I    CHNRM = 20        CHMAX = 0        RDBUF = 2091
IST579I    ATTN  = 17        TIMERS = 0        QDPTH = 14
IST580I    BUFCAP = 0        PRI = 6        SLODN = 0
IST581I    IPIU  = 19        OPIU = 20        DLRMAX = 3
IST1022I   WRBUF = 1192
IST314I    END
```

For a description of these fields, see message IST577I in [z/OS Communications Server: SNA Messages](#).

Tuning statistics file report for channel-to-channel adapters

Table 54 on page 513 shows the format of the record written to the appropriate tuning statistics file.

<i>Table 54. Record format for channel-to-channel adapters</i>			
Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	Packed	Date: 01YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	CTCA name (ID)
26	4	Binary	Dump load restart requests (DLRMAX)
30	4	Binary	Normal-sized channel program count (CHNRM)
34	4	Binary	Large-sized channel program count (CHMAX)
38	4	Binary	Total attention interrupts received (ATTN)
42	4	Binary	Number of write buffers used (WRBUF). For packed channel programs the counts represent the number of bytes read or written. For old style (non-packed format) it is the number of buffers.
46	4	Binary	Number of PIUs inbound (IPIU)
50	4	Binary	Number of PIUs outbound (OPIU)
54	4	Binary	Total input bytes used (RDBUF). For packed channel programs the counts represent the number of bytes read or written. For old style (non-packed format) it is the number of buffers.
58	4	Binary	Number of slowdowns (SLODN)
62	1	Binary	CTCA extension length (including this field)
63	1	Binary	CTCA attachment type

Table 54. Record format for channel-to-channel adapters (continued)			
Offset	Length	Format	Description
64	2	EBCDIC	CTCA version = X'F0F1'
66	4	Binary	Channel program starts because of timer trigger (TIMERS)
70	4	Binary	Channel program starts because of queue depth limit trigger (QDPH)
74	4	Binary	Channel program starts because of destination capacity limit trigger (BUFCAP)
78	4	Binary	Channel program starts because of high priority request trigger (PRI)
82	32	Binary	Reserved

Tuning I/O operations for multipath channel connections

Multipath channel (MPC) connections use XCF or channel connectivity.

Tuning I/O operations for multipath channel (MPC) connections using XCF

With XCF multipath channel attachment, VTAM uses the MVS XCF signaling facility to send data to other hosts. By analyzing statistics provided by VTAM, you can select I/O buffer sizes and other options that can improve performance of these I/O operations. The amount of data that VTAM can read in one operation depends on the size of the read buffer.

Each tuning statistics record contains information about the state of data-transfer operations between two VTAMs (using multipath channel connections). Each record contains statistics that cover the time period since the last tuning statistics record was written for that multipath channel connection.

Online tuning statistics for multipath channel adapters using XCF

Following is an example of a tuning statistics report for an XCF multipath channel attachment:

```

IST1230I TIME      = time      DATE      = date      ID = id
IST1231I IPDU      = ipdu      OPDU      = opdu
IST1569I INLP      = inlp      ONLP      = onlp
IST1232I TSWEPT    = tsweep    QSWEPT    = qsweep
IST924I  -----
IST1505I TYPE      = type      TOKEN     = token
IST1234I BSIZE     = bsize     MAXBYTES  = maxbytes
IST1236I BYTECNT0  = bytecnt0  BYTECNT   = bytecnt  DIR  = direction
IST1236I BYTECNT0  = bytecnt0  BYTECNT   = bytecnt  DIR  = direction
IST1570I NBYTECT0  = nbyTECT0  NBYTECT   = nbyTECT
IST314I  END

```

For a description of these fields, see message IST1505I in [z/OS Communications Server: SNA Messages](#).

Note: For XCF connections, *tsweep* and *qsweep* are always 0.

Tuning statistics file report for multipath channel connections using XCF

Table 55 on page 514 shows the format of the record written to the appropriate tuning statistics file. For each MPC group, two records are written. The first record contains statistics for the entire MPC group. The second record contains statistics for the READ/WRITE subchannel.

Table 55. Record format for multipath channel connections (XCF)			
Offset	Length	Format	Description
0	2	Binary	Record length

Table 55. Record format for multipath channel connections (XCF) (continued)

Offset	Length	Format	Description
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	packed	Date: 01YYDDDF Date: 00YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	TRLE name
26	4	Binary	Reserved
30	8	Binary	MVS token
38	4	Binary	Number of timer sweeps (TSWEEP) ^{1, 3}
38	4	Binary	Number of received bytes ²
42	4	Binary	Number of queue sweeps (QSWEEP) ^{1, 3}
42	4	Binary	Receive byte overflow count ²
46	4	Binary	Number of write records ¹
46	4	EBCDIC	XCF identifier (XCF) ²
50	4	Binary	Number of read records ¹
50	4	Binary	Transmit block size (BSIZE) ²
54	4	Binary	Number of READ/WRITE records ¹
54	4	EBCDIC	READ/WRITE indicator (RDWR) ²
58	4	Binary	Reserved ¹
58	4	Binary	Sent byte count ²
62	1	Binary	CTCA extension length (including this field)
63	1	Binary	CTCA attachment type X'01' ¹ X'03' ²
64	2	EBCDIC	CTCA version = X'F0F2'
66	4	Binary	Reserved ¹
66	4	Binary	Send byte overflow count ²
70	4	Binary	Reserved
74	4	Binary	Reserved
78	4	Binary	Max transmit size (MAXBYTES) ²
82	4	Binary	Reserved
86	4	Binary	Reserved
90	4	Binary	Reserved
94	4	Binary	Number of inbound NLPs (INLP)

Table 55. Record format for multipath channel connections (XCF) (continued)

Offset	Length	Format	Description
98	4	Binary	Number of outbound NLPs (ONLP)
102	4	Binary	NLP byte count
106	4	Binary	NLP byte count overflow
Notes: <ol style="list-style-type: none"> 1. Indicates fields used only in the first record for statistics on the group. 2. Indicates fields used only in the subsequent record for statistics on the read/write subchannel. 3. TSWEET and QSWEET are always 0 for HPDT MPC connections. 			

Tuning I/O operations for multipath channel connections using channels

With MPC channel connectivity attachment, VTAM performs I/O operations across multiple single-direction channel links. By analyzing statistics provided by VTAM, you can select I/O buffer sizes, data transfer delay, and other options that can improve performance of these I/O operations.

VTAM uses channel programs or Direct Memory Access (DMA) to transmit and receive data. The amount of data that VTAM can read in one operation depends on the size of the read buffer.

Each tuning statistics record contains information about the state of data-transfer operations across the multipath channel connection. Each record contains statistics that cover the time period since the last tuning statistics record was written for that multipath channel connection.

Note: Non-HPDT multipath channel uses a special data block called a sweep that is exchanged with the adjacent host to verify that data has not been lost. A host initiating a sweep request holds all outbound multipath channel transmissions until it receives a sweep reply from the adjacent host.

A sweep is initiated when either of the following occurs:

- A timer expires in the host with the higher subarea number.
- Receive queue depth in either host is excessive.

The host initiating the sweep sends the sequence number of the last output transmit block. The adjacent host compares this number with its last input transmit block sequence number. The adjacent host then sends a response to the initiating host that includes the adjacent host last output transmit block sequence number. The initiating host makes the same comparison. If the numbers do not match, or the sweep does not complete within a time limit, the multipath channel group will be deactivated. Otherwise, normal flow continues. The tuning statistics contain a count of how many sweeps are initiated by an expired timer and how many are initiated by excessive receive queue depth.

Online tuning statistics for multipath channel connections using channels

The following is an example of an MPC channel-to-channel, QDIO and Hipersockets tuning statistics report. The first nine lines are common for all three device types. The remaining lines are specific to QDIO and Hipersockets.

```

IST1230I TIME      = time      DATE      = date      ID = id
IST1231I IPDU      = ipdu      OPDU      = opdu
IST1569I INLP      = inlp      ONLP      = onlp
IST1232I TSWEET    = tsweep    QSWEEP    = qsweep
IST924I -----
IST1233I DEV       = dev       DIR        = dir
IST1234I BSIZE     = bsize     MAXBYTES  = maxbytes
IST1235I SIO       = sio       SLOWDOWN  = slowdown
IST1236I BYTECNTO  = bytecnto  BYTECNT  = bytecnt
IST1570I NBYTECTO  = nbyTECTO  NBYTECT  = nbyTECT
.
.
IST924I -----
IST1233I DEV       = dev       DIR        = dir
IST1719I PCIREALO  = pcirealo  PCIREAL  = pcireal
IST1720I PCIVIRTO  = pcivirto  PCIVIRT  = pcivirt
IST1750I PCITHRSO  = pcithrso  PCITHRSH  = pcithrsh
IST1751I PCIUNPRO  = pciunpro  PCIUNPRD  = pciunprd
IST1752I RPROCDEO  = rprocdeo  RPROCDEF  = rprocdef
IST1753I RREPLDEO  = rrepldeo  RREPLDEF  = rrepldef
IST1754I NREADSO   = nreadso   NREADS   = nreads
IST1721I SBALCNTO  = sbalcnto  SBALCNT  = sbalcnt
IST1722I PACKCNTO  = packcnto  PACKCNT  = packcnt
IST2185I FRINVCTO  = frinvcto  FRINVCT  = frinvct
IST1236I BYTECNTO  = bytecnto  BYTECNT  = bytecnt
IST1810I PKTIQDO  = pktiqdo   PKTIQD   = pktiqd
IST1811I BYTIQDO  = bytiqdo   BYTIQD   = bytiqd
IST924I -----
IST1233I DEV       = dev       DIR        = dir
IST1755I SBALMAX   = sbalmax   SBALAVG  = sbalavg
IST1756I QDPTHMAX  = qdpthmax  QDPTHAVG = qdpthavg
IST1723I SIGACNTO  = sigacnto  SIGACNT  = sigacnt
IST1721I SBALCNTO  = sbalcnto  SBALCNT  = sbalcnt
IST1722I PACKCNTO  = packcnto  PACKCNT  = packcnt
IST2242I SIGMCNTO  = sigmcnto  SIGMCNT  = sigmcnt
IST1236I BYTECNTO  = bytecnto  BYTECNT  = bytecnt
IST1810I PKTIQDO  = pktiqdo   PKTIQD   = pktiqd
IST1811I BYTIQDO  = bytiqdo   BYTIQD   = bytiqd
.
.
IST314I END

```

Figure 141. Example of an MPC channel-to-channel, QDIO, and Hipersockets tuning statistics report

For a description of these fields, see message IST1230I in [z/OS Communications Server: SNA Messages](#).

Note: For HPDT MPC connections, *tsweep* and *qsweep* are always 0.

Tuning statistics file report for multipath channel connections

Table 56 on page 517 shows the format of the record written to the appropriate tuning statistics file. For each MPC group, a minimum of three records are written. For each OSA-Express MPC group, a minimum of eight records are written. The first record contains statistics for the entire MPC group. The subsequent records contain statistics for each write subchannel, read subchannel, or OSA-Express datapath queue.

Table 56. Record format for multipath channel connections (channel)			
Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	packed	Date: 01YYDDDF Date: 00YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	MPC line name

Table 56. Record format for multipath channel connections (channel) (continued)

Offset	Length	Format	Description
26	4	Binary	Dump load restart requests (DLRMAX)
30	4	Binary	Number of inbound PIUs (IPIU) ¹
34	4	Binary	Number of outbound PIUs (OPIU) ¹
38	4	Binary	Number of timer sweeps (TSWEEP) ^{1, 3}
42	4	Binary	Number of queue sweeps (QSWEEP) ^{1, 3}
46	4	Binary	Number of write records
46	4	EBCDIC	Device address (DEV) ²
50	4	Binary	Number of read records ¹
50	4	Binary	Transmit block size (BSIZE) ²
54	4	EBCDIC	Subchannel polarity (DIR) ²
58	4	EBCDICBinary	Transmit or receive byte count (BYTECNT) ^{2, 4}
62	1	Binary	CTCA extension length (including this field)
63	1	Binary	CTCA attachment type X'01' ¹ X'02' ² X'04' ⁴
64	2	EBCDIC	CTCA version = X'F0F2'
66	4	Binary	Overflow byte count (BYTECNT0) ^{2, 4}
70	4	Binary	Slowdown frequency (SLOWDOWN) ²
74	4	Binary	Number of SIO issued (SIO) ²
78	4	Binary	Max transmit size (MAXBYTES) ²
82	4	Binary	Reserved
86	4	Binary	Reserved
90	4	Binary	Reserved
94	4	Binary	Number of inbound NLPs (INLP)
98	4	Binary	Number of outbound NLPs (ONLP)
102	4	Binary	NLP byte count
106	4	Binary	NLP byte count overflow
110	4	Binary	Number of OSA-Express datapath queues ¹
114	4	Binary	Reserved
118	4	Binary	OSA-Express PCIR overflow for READ queue ⁴ OSA-Express maximum SBALs for WR/x queue ⁴
122	4	Binary	OSA-Express PCIR count for READ queue ⁴ OSA-Express average SBALs for WR/x queue ⁴

Table 56. Record format for multipath channel connections (channel) (continued)

Offset	Length	Format	Description
126	4	Binary	OSA-Express PCIV overflow for READ queue ⁴ OSA-Express maximum queue depth for WR/x queue ⁴
130	4	Binary	OSA-Express PCIV count for READ queue ⁴ OSA-Express average queue depth for WR/x queue ⁴
134	4	Binary	OSA-Express PCIT overflow for READ queue ⁴ OSA-Express SIGA overflow for WR/x queue ⁴
138	4	Binary	OSA-Express PCIT count for READ queue ⁴ OSA-Express SIGA count for WR/x queue ⁴
142	4	Binary	OSA-Express PCIU overflow for READ queue ⁴ OSA-Express SBAL count overflow for WR/x queue ⁴
146	4	Binary	OSA-Express PCIU count for READ queue ⁴ OSA-Express SBAL count for WR/x queue ⁴
150	4	Binary	OSA-Express processing deferrals overflow for READ queue ⁴ OSA-Express packet count overflow for WR/x queue ⁴
154	4	Binary	OSA-Express processing deferrals for READ queue ⁴ OSA-Express packet count for WR/x queue ⁴
158	4	Binary	OSA-Express replenishment deferrals overflow for READ queue ⁴ OSA-Express or Hipersockets accelerated packet count overflow for WR/x queue ⁴
162	4	Binary	OSA-Express replenishment deferrals for READ queue ⁴ OSA-Express or Hipersockets accelerated packet count for WR/x queue ⁴
166	4	Binary	OSA-Express reads exhausted overflow for READ queue ⁴ OSA-Express or Hipersockets accelerated byte count overflow for WR/x queue ⁴
170	4	Binary	OSA-Express reads exhausted for READ queue ⁴ OSA-Express or Hipersockets accelerated byte count for WR/x queue ⁴

Table 56. Record format for multipath channel connections (channel) (continued)

Offset	Length	Format	Description
174	4	Binary	OSA-Express SBAL count overflow for READ queue ⁴ Undefined for WR/x queue ⁴
178	4	Binary	OSA-Express SBAL count for READ queue ⁴ Undefined for WR/x queue ⁴
182	4	Binary	OSA-Express early interrupt count overflow for READ queue ⁴ Undefined for WR/x queue ⁴
186	4	Binary	OSA-Express early interrupt count for READ queue ⁴ Undefined for WR/x queue ⁴
190	4	Binary	OSA-Express early interrupt phase II spinout count overflow for READ queue ⁴ Undefined for WR/x queue ⁴
194	4	Binary	OSA-Express early interrupt phase II spinout count overflow for READ queue ⁴ Undefined for WR/x queue ⁴
198	4	Binary	OSA-Express packet count overflow for READ queue ⁴ Undefined for WR/x queue ⁴
202	4	Binary	OSA-Express packet count for READ queue ⁴ Undefined for WR/x queue ⁴
206	4	Binary	OSA-Express or Hipersockets accelerated packet count overflow for READ queue ⁴ Undefined for WR/x queue ⁴
210	4	Binary	OSA-Express or Hipersockets accelerated packet count for READ queue ⁴ Undefined for WR/x queue ⁴
214	4	Binary	OSA-Express or Hipersockets accelerated byte count overflow for READ queue ⁴ Undefined for WR/x queue ⁴
218	4	Binary	OSA-Express or Hipersockets accelerated byte count for READ queue ⁴ Undefined for WR/x queue ⁴
222	4	Binary	OSA-Express3 or later frame invalidation packet count overflow for READ queue ⁴ Undefined for WR/x queue ⁴

Table 56. Record format for multipath channel connections (channel) (continued)

Offset	Length	Format	Description
226	4	Binary	OSA-Express3 or later frame invalidation packet count for READ queue ⁴ Undefined for WR/x queue ⁴
258	32	Binary	Reserved
262	8	EBCDIC	READ queue name ⁴ Undefined for WR/x queue ⁴

Notes:

1. Indicates fields used only in the first record for statistics on the group.
2. Indicates fields used only in subsequent records for statistics on each write and each read subchannel.
3. TSWEET and QSWEEP are always 0 for HPDT MPC connections.
4. Indicates fields used only in subsequent records for statistics on each OSA-Express or Hipersockets datapath queue.

Tuning input/output (I/O) operations for TCP connections

VTAM performs I/O operations across TCP connections. By analyzing statistics provided by VTAM, you can improve performance of I/O operations.

Online tuning statistics report for TCP connections

Following are examples of tuning statistics reports for TCP connections. For a description of these fields, see the appropriate messages in [z/OS Communications Server: SNA Messages](#).

CDLC

```

IST1230I TIME      = 16500936      DATE      = 01017      ID = IUTC00B2
IST1613I TYPE      = CDLC          ATTN      =          38
IST1653I RWSIO     =          83    WCH       =          46    RCH =          43
IST1654I INPACKET  =          50    INBYTE    =          5516  MAX =          243
IST1655I OTPACKET  =          51    OTBYTE    =          5841  MAX =          255
IST314I END

```

CTC

```

IST1230I TIME      = 17050069      DATE      = 01017      ID = IUTX0514
IST1613I TYPE      = CTC          ATTN      =          0
IST1614I READSIO   =          43    PACKET    =          48    BYT =          5594
IST1615I ARPACKEt  =          1    ARBYTE    =          130    MAX =          249
IST1616I WRITESIO  =          44    PACKET    =          49    BYT =          5965
IST1617I AWPACKEt  =          1    AWBYTE    =          135    MAX =          241
IST314I END

```

HyperChannel

```

IST1230I TIME      = 17102276      DATE      = 01017      ID = IUTH0546
IST1613I TYPE      = HYP          ATTN      =          3
IST1614I READSIO   =          53    PACKET    =          53    BYT =          6529
IST1615I ARPACKEt  =          1    ARBYTE    =          123    MAX =          259
IST1616I WRITESIO  =          53    PACKET    =          53    OUT =          6791
IST1617I AWPACKEt  =          1    AWBYTE    =          128    MAX =          271
IST314I END

```

Tuning statistics file report for TCP connections

If statistics are produced, they are always written to the appropriate tuning statistics file in the format shown in [Table 57](#) on page 522.

<i>Table 57. Record format for TCP connections</i>			
Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	Packed	Date: 00YYDDDF Date: 01YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	TCP line name
26	4	Binary	Reserved
30	4	Binary	Count of write channel programs (not CDLC)
34	4	Binary	Count of read channel programs (not CDLC)
38	4	Binary	Number of attentions
42	4	Binary	Largest outbound packet sent
46	4	Binary	Largest packet received
50	4	Binary	Reserved
54	4	Binary	Reserved
58	4	Binary	Reserved
62	1	Binary	Extension length (including this field)
63	1	Binary	Reserved
64	2	EBCDIC	TCP version = X'F0F3'
66	4	Binary	Reserved
70	4	Binary	Reserved
74	4	Binary	Reserved
78	4	Binary	Reserved
82	4	Binary	Inbound packet count
86	4	Binary	Outbound packet count
90	4	Binary	Inbound byte count
94	4	Binary	Inbound byte count, overflow
98	4	Binary	Outbound byte count
102	4	Binary	Outbound byte count, overflow

Table 57. Record format for TCP connections (continued)

Offset	Length	Format	Description
106	1	Binary	TCP legacy type <ul style="list-style-type: none"> • X'10' = CTC • X'40' = CDLC • X'50' = HYPERchannel • X'60' = SameHost
107	3	Binary	Reserved
110	4	Binary	Number of PCI interrupts (CLAW only)
114	4	Binary	Number of READ CCWs completed (CLAW only)
118	4	Binary	Number of WRITE CCWs completed (CLAW only)
122	4	Binary	Number of WRITE appends (CLAW only)
126	4	Binary	Number of READ/WRITE SIOs (CLAW only)

Tuning input/output (I/O) operations for RoCE connections

Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) uses "RoCE Express" features to enable read/write access between the memories of connected hosts.

VTAM maintains tuning statistics for each "RoCE Express" port that you can display at VTAM tuning statistics intervals. Use the MODIFY TNSTAT,TRLE=RNIC_TRLEname command to enable or disable the statistics.

All statistics are cumulative totals reported as 8-byte values with a high-word overflow. The user-level statistics are also reported by NMI.

Online tuning statistics report for RoCE connections

This is an example of a tuning statistics report for a "RoCE Express" port and its users. The message output for each "RoCE Express" port includes the global "RoCE Express" statistics, followed by sections for inbound and outbound statistics for each user. For a description of the fields, see the appropriate messages in [z/OS Communications Server: SNA Messages](#).

```

IST1230I TIME      = 18051835   DATE      = 09182       ID   = EZARIUT1FEFC
IST1719I PCIREALO  = 0          PCIREAL   = 5
IST1751I PCIUNPRO  = 0          PCIUNPRD  = 3
IST2366I POLLEQO   = 0          POLLEQ    = 15
IST2367I POLLEQEO  = 0          POLLEQE   = 25
IST924I -----
IST2368I ULP_ID    = TCPCS1
IST2369I POLLCQO   = 0          POLLCQ     = 250
IST2370I POLLCQUO  = 0          POLLCQU    = 50
IST2371I POLLCQEO  = 0          POLLCQE    = 1800
IST2372I SRBSCHDO  = 0          SRBSCHD    = 15
IST2373I SRBRSCO   = 0          SRBRSCHD   = 0
IST2374I INBBYTLO  = 0          INBBYTEL   = 176
IST2375I INBBYTMO  = 0          INBBYTEM   = 7200
IST2376I INBBYTNO  = 0          INBBYTEN   = 305306
IST2377I DATAREQO  = 0          DATAREQ    = 60
IST2378I POSTO     = 0          POST       = 70
IST2379I POSTEO    = 0          POSTELEM   = 178
IST2380I POSTQEO   = 0          POSTQUED   = 10
IST2381I OUTBYTLO  = 0          OUTBYTEL   = 176
IST2382I OUTBYTMO  = 0          OUTBYTEM   = 4000
IST2383I OUTBYTNO  = 0          OUTBYTEN   = 58950

```

Message subgroup IST2368I through IST2383I is repeated for each TCP/IP stack that activated the "RoCE Express" port.

Tuning statistics file report for RoCE connections

Statistics that are produced are always written to the appropriate tuning statistics file in the format shown in Table 58 on page 524. For each RoCE connection, a minimum of two records are written; the first record contains global statistics for the RoCE port, and the subsequent records contain statistics for an individual RoCE user.

<i>Table 58. Record format for RoCE connections</i>			
Offset	Length	Format	Description
0	2	Binary	Record length
2	3	Binary	Reserved
5	1	Binary	Record type
6	4	Binary	Time record moved to buffer
10	4	Packed	Date: 01YYDDDF
14	4	EBCDIC	System identification
18	8	EBCDIC	RoCE connection name
26	4	Binary	Reserved
30	4	Binary	Reserved
34	4	Binary	Reserved
38	4	Binary	Reserved
42	4	Binary	Reserved
46	4	Binary	Reserved
50	4	Binary	Reserved
54	4	Binary	Reserved
58	4	Binary	Reserved
62	1	Binary	Extension length (including this field)
63	1	Binary	Attachment type X'01' RoCE port X'02' RoCE user
64	2	EBCDIC	Version = X'F0F5'
66	4	Binary	Reserved
70	4	Binary	Reserved
74	4	Binary	Reserved
78	4	Binary	Reserved
82	4	Binary	Number of RoCE users “1” on page 526
86	4	Binary	Poll EQ overflow “1” on page 526
90	4	Binary	Poll EQ count “1” on page 526

Table 58. Record format for RoCE connections (continued)

Offset	Length	Format	Description
94	4	Binary	Poll EQ entries overflow “1” on page 526
98	4	Binary	Poll EQ entries count “1” on page 526
102	4	Binary	PCI real interrupts overflow “1” on page 526
106	4	Binary	PCI real interrupts count “1” on page 526
110	4	Binary	Unproductive PCI overflow “1” on page 526
114	4	Binary	Unproductive PCI count “1” on page 526
118	8	EBCDIC	ULP ID “2” on page 526
126	4	Binary	Outbound data request overflow “2” on page 526
130	4	Binary	Outbound data request count “2” on page 526
134	4	Binary	Outbound RDMA request overflow “2” on page 526
138	4	Binary	Outbound RDMA request count “2” on page 526
142	4	Binary	Outbound RDMA operation overflow “2” on page 526
146	4	Binary	Outbound RDMA operation count “2” on page 526
150	4	Binary	Queued outbound RDMA request overflow “2” on page 526
154	4	Binary	Queued outbound RDMA request count “2” on page 526
158	4	Binary	Outbound RoCE inline data bytes overflow “2” on page 526
162	4	Binary	Outbound RoCE inline data bytes count “2” on page 526
166	4	Binary	Outbound RoCE immediate data bytes overflow “2” on page 526
170	4	Binary	Outbound RoCE immediate data bytes count “2” on page 526
174	4	Binary	Outbound RDMA data bytes overflow “2” on page 526
178	4	Binary	Outbound RDMA data bytes count “2” on page 526
182	4	Binary	RoCE inbound poll request overflow “2” on page 526
186	4	Binary	RoCE inbound poll request count “2” on page 526
190	4	Binary	RoCE inbound poll request operation overflow “2” on page 526

Table 58. Record format for RoCE connections (continued)

Offset	Length	Format	Description
194	4	Binary	RoCE inbound poll request operation count ^{“2”} on page 526
198	4	Binary	Unproductive RoCE inbound poll request overflow ^{“2”} on page 526
202	4	Binary	Unproductive RoCE inbound poll request count ^{“2”} on page 526
206	4	Binary	Interrupt handler SRB dispatch overflow ^{“2”} on page 526
210	4	Binary	Interrupt handler SRB dispatch count ^{“2”} on page 526
214	4	Binary	Additional SRB dispatch overflow ^{“2”} on page 526
218	4	Binary	Additional SRB dispatch count ^{“2”} on page 526
222	4	Binary	Inbound RoCE inline data bytes overflow ^{“2”} on page 526
226	4	Binary	Inbound RoCE inline data bytes count ^{“2”} on page 526
230	4	Binary	Inbound RoCE immediate data bytes overflow ^{“2”} on page 526
234	4	Binary	Inbound RoCE immediate data bytes count ^{“2”} on page 526
238	4	Binary	Inbound RDMA data bytes overflow ^{“2”} on page 526
242	4	Binary	Inbound RDMA data bytes count ^{“2”} on page 526
Note: 1. Indicates fields used only in the first record for statistics on the RoCE port. 2. Indicates fields used only in subsequent records for statistics on each RoCE user.			

Analyzing tuning statistics

A single set of VTAM tuning statistics can be enough to indicate how a network is operating. However, these statistics become more valuable as you compare sets of values over time to see trends or the effects caused by changing buffer pool specifications and operands. You should analyze tuning statistics before and after making any change that might affect system performance.

Use tuning statistics to analyze the effects of VTAM host processor usage.

1. Compare the number of attention interrupts with channel READs. If the number of attention interrupts is the same as the number of channel READs, each READ is the result of the NCP sending the host a stand-alone attention. This condition causes increased CPU utilization because the host processor and VTAM must handle each attention interrupt. Conversely, if the number of attention interrupts is small in relation to the number of channel READs, the impact on the host is much less. If the number of attentions is large in relation to the number of channel READs (CHRD or RCH), then this could be a normal condition. Some controllers (for example, those running NCP Version 5 Release 2 and later

releases) send attention polls to the primary logical unit during periods of no traffic. VTAM does not count the number of channel reads when no data is transferred.

2. Analyze your IOBUF size. If RDBUF is much larger than IPIU, the value specified for the I/O buffer size might be too small. However, you need to know the message characteristics of your network before you make decisions. Many systems receive the bulk of their inbound data in very short messages and respond with longer outbound messages. For details on choosing an appropriate value for the I/O buffer size, see [“Guidelines for setting UNITSZ”](#) on page 553.
3. Analyze data transfer operations between VTAM and an SNA controller. To analyze inbound data transfer, multiply MAXBFRU by the number of channel READs (CHRD or RCHD). MAXBFRU is specified on the GROUP, HOST, LINE, or PU definition statement. The product should be close to the value of RDBUF (the number of read buffers used). Otherwise, data is not being transferred for all the channel command words (CCWs) in the read channel program. Reducing the value of MAXBFRU should solve this problem. However, the value of MAXBFRU times UNITSZ must be at least as large as the largest PIU received. For more details on setting the value of MAXBFRU and UNITSZ, see [“Guidelines for setting UNITSZ”](#) on page 553 and [“Setting the MAXBFRU operand”](#) on page 551.
4. Examine the read attention (RDATN) information that is available in VTAM tuning statistics to understand the efficiency of the inbound data transfer operations. RDATN is the number of times that VTAM, after reading data, is requested by an attention to read more data. This happens when one of the following occurs:
 - The READ channel program CCW string is not long enough to contain all of the data sent.
 - Enough additional data comes into the communication controller during the read operation to cause it to request VTAM to do another read.
 - The controller is short on buffer space and is requesting a read to confirm VTAM receipt of data sent by the controller, so that the controller buffers can be reused.

Although this coattailing of PIUs is preferable to stand-alone attention interrupts, a large RDATN value is not desirable. To increase the size of the CCW string, increase the MAXBFRU value. For more details on setting the value of MAXBFRU, see [“Setting the MAXBFRU operand”](#) on page 551.

5. Analyze the outbound data transfer operation between VTAM and an SNA controller by analyzing VTAM blocking of outbound PIUs. To determine the average number of PIUs for each write operation, divide the outbound PIU (OPIU) count by the number of channel WRITES (CHWR or WCH). This average number of PIUs indicates the effectiveness of the VTAM blocking algorithm. Increasing DELAY can also increase inbound coattailing. For more information about setting the value of DELAY, see [“Guidelines for setting DELAY”](#) on page 549.

Determining the amount of coattailing in your system

You can analyze the amount of coattailing in your system by examining VTAM tuning statistics. The tuning statistics indicate the total number of READ (CHRD or RCH) and WRITE (CHWR or WCH) channel programs that are used for data transfer operations. This information can be used to tune your system to improve coattailing if necessary. Increasing coattailing in your system can impact response times because coattailing means there are data transfer operations delays. When examining tuning statistics, also consider that short periods of high-attention counts can indicate a light load on the network.

Migrating from user-replaceable constants

User-replaceable constants were formerly used to control various VTAM functions, such as:

- Buffer use storage management services (SMS) trace “snapshot” value
- Maximum RU size for sessions
- BSC 3270 timeout (inoperative) value
- Virtual-route-selection subtask interval analysis block
- Host intermediate routing node (IRN) slowdown storage
- NetView trace data buffers

- Directory size of the symbol resolution table for the network

User-replaceable constants are now set only as start options. For additional information, see [z/OS Communications Server: SNA Customization](#). For more information about start options, see the [z/OS Communications Server: SNA Resource Definition Reference](#). Start options can be displayed with the DISPLAY VTAMOPTS command, and many of them can be changed with the MODIFY VTAMOPTS command. For more information about the DISPLAY VTAMOPTS and MODIFY VTAMOPTS commands, see [z/OS Communications Server: SNA Operation](#).

Estimating the number of active sessions

The EAS operand on the LU and APPL definition statements specifies the estimated number of active sessions between VTAM and the application program or VTAM and the logical unit.

EAS operand for application programs

The EAS operand on the APPL definition statement indicates the amount of storage to allocate for a table that contains information about any sessions with the application program. The EAS value that you choose can affect performance in your system.

Specifying an EAS value that is too small can cause an increase in path length for VTAM processing of inbound data flows. The increase is because of VTAM scanning control blocks for some of the sessions. As more sessions become active, the directory entries are used more frequently, and the chains become longer. Longer chains require more scan time. If you specify a larger EAS value, the chains for particular directory entries may be smaller and may use less CPU time for scanning the chains.

For a non-TSO application program, you should specify an EAS value that equals the average number of sessions, but does not exceed 32 767. If you have virtual storage constraints, you can specify a smaller EAS value.

For a TSO application program, you should specify EAS=1.

EAS operand for independent logical units

For performance reasons, if you anticipate that a particular logical unit can have more than 256 sessions, increase the value of the EAS operand to the nearest multiple of 256. For example, if you anticipate 700 sessions for a particular logical unit, you should define EAS to be 768 (3 x 256). Coding a poor estimate for EAS does not affect your sessions, but it does affect performance.

You can define the EAS operand for an independent logical unit only.

Common storage areas

Common storage areas (CSA) are used by VTAM and other resources in your network to maintain buffers and control blocks in the CSA for MVS systems. For performance reasons, VTAM offers you some control over its use of common storage areas with the start options. These options are described in the following sections.

Common service area limit

The CSALIMIT start option specifies the maximum amount of common service area (CSA) that is used by VTAM. It can be used to prevent VTAM from using CSA that is needed by the operating system. Most environments do not need to limit VTAM use of CSA because their peak CSA usage is a very small increment over their requirements when all of their logical units are in session.

If the limit specified by CSALIMIT is reached, the results are unpredictable. If LPBUF cannot be expanded, VTAM might enter a deadlock condition. Therefore, if you specify the CSALIMIT start option, you should define LPBUF large enough so that it does not have to expand. Other possible occurrences are that messages can be lost or a session initiation or termination can fail. Coding a value for CSALIMIT without

the ,F modifier can limit these exposures. See the [z/OS Communications Server: SNA Resource Definition Reference](#) for details.

Common service area 24-bit

The CSA24 start option specifies the maximum amount of 24-bit addressable common service area (CSA) that can be used by VTAM. The largest possible value for CSA24 is 16 megabytes.

The CSA24 start option limits only the amount of 24-bit addressable storage explicitly requested by VTAM. The CSA24 limitations do not apply when MVS returns 24-bit addressable storage in response to a VTAM storage request above the 16-megabyte line. MVS returns 24-bit addressable storage below the line because not enough CSA above the line is defined. CSA storage can be allocated from above or below the line depending on where the allocating procedure resides.

DISPLAY STORUSE pools

The DISPLAY STORUSE command provides a way to remedy a possible shortage of storage space. Table 59 on page 529 provides a list of storage pools that are displayed using the DISPLAY STORUSE command. Included for each pool is a short description of the pool function and characteristics. These pools are not customer-defined, unlike the buffer pools defined using the VTAM start options (for example, IOBUF). VTAM allocates and deallocates storage from these pools as needed.

If VTAM is in a storage shortage situation, Table 59 on page 529 and the output from the DISPLAY STORUSE command can be used to determine where excess storage is being used, enabling you to take appropriate action to remedy the shortage.

Table 59. DISPLAY STORUSE pools		
Pool name	Storage location	Description
ACDEB	SYSTEM	A pool element is allocated for every active application.
ACPCB	USER	An element is allocated for every adjacent control point with which this node has an active CP-CP session.
ADJCP	USER	Each pool element defines a single adjacent control point (ADJCP).
ADJNODE	USER	Elements are allocated for each CP-CP session partner to track topology flow status with an adjacent node.
ALPHCD	SYSTEM	Element used by HPDT MPC to manage CSM buffer descriptors. These descriptors represent the storage used for the physical transmission of data over HPDT MPC.
AMU	SYSTEM	Elements are short-lived signals used for intraproduct communication.
ANDCB	USER	One element is allocated the first time a node activates a link supporting CP-CP sessions with a specific adjacent node.
ATGB	USER	An element represents a single T2.1 connection or VR-TG to an adjacent CP.
AUTOLOGN	USER	Elements are used to keep track of autologon relationships.
BFRTRACE	SYSTEM	Elements are used to hold small buffer trace records.
BFRTRFUL	SYSTEM	Elements are used to hold large buffer trace records.
BSBEXT	USER	Elements are allocated for each session using SNA/IP support.
CAB	SYSTEM	One element is allocated for each VCNS connection resulting from a LOGON to a VCNS line (LANs) or from an X.25 OPEN command.

Table 59. DISPLAY STORUSE pools (continued)

Pool name	Storage location	Description
CACHE	USER	Elements are used for caching PCIDs during direct search list processing.
CANT	SYSTEM	One element is allocated for every 64 VCNS X.25 connections for the same LOGON to a VCNS line.
CDAJSCP	USER	Elements are used to define adjacent SSCP entries.
CDRSC	USER	Elements are used to define dynamic CDRSCs and clone CDRSCs.
CFSACCCD	SYSTEM	Coupling facility short-lived common storage pool.
CFSACCCS	SYSTEM	Coupling facility long-lived common storage pool.
CFSACCPD	USER	Coupling facility short-lived private storage pool.
CFSACCPD	USER	Coupling facility long-lived private storage pool.
CFSBUFCD	SYSTEM	Elements are short-lived buffer objects used to manage coupling facility structure data buffers.
CFSBUFCS	SYSTEM	Elements are long-lived buffer objects used to manage coupling facility structure data buffers.
CFSBUFPD	USER	Elements are short-lived buffer objects used to manage coupling facility structure.
CFSBUFPS	USER	Elements are long-lived buffer objects used to manage coupling facility structure.
CFSCSA	SYSTEM	Coupling facility common storage pool.
CFSPRIV	USER	Coupling facility private storage pool.
CMOBJ	SYSTEM	Elements are used by VTAM connection manager to represent logical connections using HPDT DLCs.
CNSFACUD	SYSTEM	An element is allocated every time a VCNS application asks VTAM to initiate or receive an X.25 call request that will contain facilities or call user data.
CORCB	USER	Elements are used for APPN Locate request correlation.
COS	USER	Elements are used for APPN COS definitions and mode table mappings.
CPRUPE	USER	Elements are request/response unit processing elements used when processing APPN-related requests.
CPWACSA	SYSTEM	An element is allocated when a USS command is entered from the network operator console or from a user terminal. One element from the pool is allocated when an application resource definition specifying SSCPFM=USSNOP is being processed.
CPWAPVT	USER	One element is allocated when a USS command is entered from the network operator console or from a user terminal. One element from the pool is allocated when an application resource definition specifying SSCPFM=USSNOP is being processed.
DCX	SYSTEM	Pool elements are used to maintain data compression information.

Table 59. DISPLAY STORUSE pools (continued)

Pool name	Storage location	Description
DDEL	USER	Elements are used to delay the disconnection of a PU that is defined with the DISCNT=DELAY parameter.
DECB	USER	An element is allocated for each resource in the APPN directory database.
DISKIO	USER	One element is allocated per component performing database hardening. The storage for this pool is allocated below the 16-M line.
DMTSQ	USER	One element is allocated every time a message contained in the message flooding table is issued. The allocated element is freed when the message suppression time expires.
DSERVER	USER	Elements are control blocks and short-term signals related to directory services and interchange nodes.
DSSIB	USER	An element is allocated when a DSRLST is received and freed when the DSRLIST response is sent.
DSUTIL	USER	Elements are used to perform locate search processing.
DYPATH	USER	Elements are path table entries for dynamic PUs.
EEHNMIPD	SYSTEM	An element from this pool is allocated when the SNAMGMT start option is set to YES and a client application connects to the SNA Network Management socket. More elements are allocated when a response is built because of a request from a client application. These additional elements are freed when the response is returned. The original element is freed when either the client or VTAM terminates the connection.
EPTDVT	SYSTEM	Elements are used to contain DLC-specific information.
ERICPOOL	USER	Elements are used while parsing request/response units. When the RU processing completes, the elements are freed.
ERTE	USER	Elements are used to define explicit routes resulting from PATH statements.
FMCB	SYSTEM	Elements are allocated once for each session associated with an application LU (for nonpersistent LU sessions) at session BIND time.
FMCBEXT	SYSTEM	Elements are allocated once for each session associated with an application LU at session BIND time.
FMH5	SYSTEM	One element is allocated for each incoming LU 6.2 conversation request.
GRINS	USER	Elements are used to maintain associations between network resources and generic names.
GWNAJSCP	USER	Elements are: <ul style="list-style-type: none"> • Adjacent SSCP routing tables used to route CDINIT or DSRLST RUs. Elements are freed when the routing completes or fails. • Information to determine the gateway NCP to use during session setup.
HIPOOLPS	USER	HPR table used by MNPS

Table 59. DISPLAY STORUSE pools (continued)

Pool name	Storage location	Description
HPRINFO	USER	Elements exist for each RTP connection for which data is being collected by a single performance monitor application.
HSICB	SYSTEM	Elements are allocated once for each APPC session associated with an application LU at session BIND time.
HSQH	SYSTEM	Elements are allocated once for the first of every 107 sessions set up across a VR, at session BIND time.
IOBLOCKL	SYSTEM	Elements in the pool are large DLC-related control blocks (for example a channel-attached NCP). The storage for this pool is fixed and is not paged out of memory by the operating system.
IOBLOCKP	USER	Elements in the pool are large DLC-related control blocks. The storage for this pool is not fixed and can be paged out of memory by the operating system.
IOBLOCKS	SYSTEM	Elements in the pool are small DLC-related control blocks (such as for a channel-attached NCP). The storage for this pool is fixed and is not paged out of memory by the operating system.
IOSIB	USER	Elements are used to process Init_Other (Cross-domain) requests.
IPADDR	USER	Element used to store IP addresses and host names.
ISTENDEL	USER	One element is allocated by a network node for each adjacent served end node or nonnative network node.
ISTSITCB	USER	One element is needed at endpoint and network node server roles for each APPN Search procedure.
ISTTRCEL	USER	Elements are used for the problem determination (PD) trace function of the CNM interface.
KEYTOKEN	SYSTEM	Elements map active cryptographic session key tokens.
LCB	USER	An element is allocated to process locate search requests.
LMTABLE	SYSTEM	Three types of elements come from this pool: <ul style="list-style-type: none"> • One element is allocated for every pair of LUs that have negotiated APPC session limits. The element is allocated when a CNOS with a partner LU is initiated, and freed when the application closes its ACB. • An element is used for every application that opens its ACB with APPC=YES. • Elements represents current session limits between two LU 6.2s on a particular mode. An element is allocated for every logmode that has been negotiated between two partner LUs. The elements are freed for deletion, or freed when the application closes its ACB.
MRPOOLPS	USER	Elements contain MNPS RTP information.
NDREC	USER	Elements are used for APPN topology node information.
NIDCB	USER	An element is allocated for each network identifier known to the APPN directory database.

Table 59. DISPLAY STORUSE pools (continued)

Pool name	Storage location	Description
NLPDELPD	USER	Elements contain MNPS NLP entry IDs.
NQDAT	USER	One element is allocated for each network-qualified SRTE.
NSRUL	SYSTEM	Elements are used to process LU 6.2 session activation. NSRUL is used for larger-sized requests.
NSRUS	SYSTEM	Elements are used to process LU 6.2 session activation. NSRUS is used for smaller-sized requests.
NSS	SYSTEM	Elements are used to process LU 6.2 session activation.
OSCB	USER	Elements are used to track outstanding locate search requests.
PAGBLBSB	USER	Elements are allocated for each session using HPR or SNA/IP support.
PAQ	USER	Elements are used to track PLU network addresses for a given LU.
PCDCA	USER	Elements are used for border node PCID caching.
PGIOBLK	SYSTEM	Elements are used when communicating with the 3172 or the OSA.
PLOCB	USER	Elements are used to process locate search requests and replies.
PLUSC	SYSTEM	Elements are used during persistent LU session CLOSE processing.
PLUSDATA	SYSTEM	Elements are allocated once for each persistent LU session associated with an application LU, at session BIND time. An additional element is allocated for each MNPS session at session BIND time.
PLUSFMCB	SYSTEM	Elements are allocated once for each persistent LU session associated with an application LU, at session BIND time.
POAPRIV	USER	One element is allocated for each message destined for a program operator application (POA). If the message requires a reply, a second POAPRIV element is allocated.
POWECOMM	SYSTEM	One element is allocated for every message issued when VTAM is running under a user task.
POWEPRIV	USER	One element is allocated for every message issued when VTAM is running under the VTAM task.
POWMCOMM	SYSTEM	One element is allocated for every single-line message and for every message group when VTAM is running under a user task.
POWMPRIV	USER	One element is allocated for every single-line message and for every message group when VTAM is running under the VTAM task.
PRDLE	SYSTEM	Elements represent random data being used for establishing LU 6.2 sessions with session-level security.
PRIDBLK	USER	Elements map procedure-relation identifier blocks (PRIDs)
PRIDQAB	USER	Elements are used to maintain procedure-relation identifier blocks (PRIDs).
PULURDTE	USER	Elements are used to define dynamic and predefined PUs and LUs.

Table 59. DISPLAY STORUSE pools (continued)

Pool name	Storage location	Description
PVTSTATC	USER	Generic utility pool for large blocks of storage that must be in VTAM private storage for an extended period of time.
PXBFIXED	SYSTEM	Elements are used to expand fixed buffer pools. The storage for this pool is fixed and is not paged out of memory by the operating system.
PXBPAGED	SYSTEM	Elements are used to expand pageable buffer pools.
RAB	SYSTEM	One element is allocated for each APPC conversation.
RAQ	USER	Elements are used to queue requests when a usable network address is not available and an RNAA must be sent.
RIBRANT	SYSTEM	One element is allocated for each LOGON of a VCNS application to a VCNS line, plus one extra element for every 16 LOGONs.
RPMNPSPS	USER	RTP elements used by MNPS
RTPINFO	USER	Contains a number of different elements, all of which are used for HPR. Both the RTP and RCM components use this pool.
RUCON	USER	One element is allocated each time a DISPLAY ROUTE,TEST=YES command is issued. After 168 elements have been allocated, VTAM will delete any elements that have not been used for more than 30 minutes.
RUPECOMM	SYSTEM	Elements are allocated from this pool to process request/response units (RUs) when execution is taking place in the VTAM address space.
RUPEPRIV	USER	Elements are allocated from this pool to process request/response units (RUs) when execution is taking place in a non-VTAM address space.
SAB	SYSTEM	One element is allocated for each APPC session.
SCCB	USER	Elements are used to perform search concentration.
SIB	USER	One element is allocated for each LU-LU session.
SIBEXT	USER	One or two elements are allocated for each cross-network session. The elements are freed when the session ends.
SIBIX	USER	One element is allocated during session initiation. The element is freed when the session becomes active.
SLD	SYSTEM	Elements are allocated to process APPCCMD CONTROL=OPRCNTL,QUALIFY=DISPLAY macro instructions, including those issued from the operator console.
SLENT	USER	A pool element is allocated when a CPSVCMG session is activated between an end node and serving network node, or for any CP SNASVCMG sessions that VTAM management services transport activates.
SM3270	HVComm	Primarily used to contain 3270 screen maps. CSM HVCOMM is used for this pool.
SPTPOOL	SYSTEM	Holds all of the SPTAEs for the associated pools.

Table 59. DISPLAY STORUSE pools (continued)

Pool name	Storage location	Description
SRTE	USER	Elements are entries in the symbol resolution table.
SSCPFMCB	USER	One element is allocated for each SSCP-PU or SSCP-LU session.
STB	USER	Elements store information concerning a T2.1 adjacent link station that an independent LU is using for session connectivity.
TCPIOCD	SYSTEM	TCP/IP IO buffer pool used for QDIO.
TGP	USER	One element is created for each TG profile (TGP) entry.
TGREC	USER	Elements are used for APPN topology TG information.
TIPACX	SYSTEM	Elements contain control information to support HPDT services. The storage for this pool is fixed and is not paged out of memory by the operating system.
TREEBLD	USER	Elements are used for APPN routing tree construction and maintenance.
TRSINFO	USER	Elements are used for topology broadcast lists (for use during topology database update broadcasting) and endpoint TG vector information (during route calculation).
UECB	SYSTEM	One element is allocated each time a user exit is to be scheduled.
UNSOL	USER	An element is allocated for every adjacent control point with which this node has an active CP-CP session.
UTILCSAL	SYSTEM	Generic utility pool for large blocks of storage that must be in CSA.
UTILCSAS	SYSTEM	Generic utility pool for small blocks of storage that must be in CSA.
UTILPRTL	USER	Generic utility pool for large blocks of storage that must be in private storage.
UTILPVTS	USER	Generic utility pool for small blocks of storage that must be in private storage.
UVRPL	USER	One element is allocated each time a user exit is to be scheduled.
VRDCB	SYSTEM	An element exists for each virtual route for which data is being collected by at least one performance monitor application.
VRPL	SYSTEM	Elements are VTAM copy of an application request parameter list (RPL). Elements are also used for BINDs and other RUs received from the network.
VRRSB	USER	An element exists for each virtual route for which data is being collected by a single performance monitor application.
WAR	USER	Elements represent autologon session originators (OLU-SLU) waiting for the availability of a required PU resource.
WREEID	USER	Elements are used to suspend and resume VTAM processes.
XNINFO	USER	Elements are used for search processing.

Table 59. DISPLAY STORUSE pools (continued)

Pool name	Storage location	Description
<p>Note: Unless otherwise specified:</p> <ul style="list-style-type: none"> • All pool storage is pageable and can be paged out of memory by the operating system. • All pool storage can be located above or below the 16M line. • Some of the pools above are defined to be associated with certain VTAM tasks. Pools are associated with VTAM tasks to improve performance during storage allocation. The DISPLAY STORUSE command will not display usage for the associated pools. The total storage used does account for storage allocated by these pools 		

Buffer pools

VTAM uses buffer pools to control the handling of data. It dynamically allocates and deallocates space in buffer pools for the VTAM control blocks, I/O buffers, and channel programs that control the transmission of data. Specifying large buffer pools can waste storage, but does not require frequent CPU use for expansion. Specifying small buffer pools conserves storage, but requires frequent CPU use for expansion and contraction. With proper tuning, you can achieve a balance between storage use and performance that is suitable for your environment.

Types of buffer pools

Table 60 on page 536 lists the types of buffer pools.

Fixed buffer pools are allocated in extended CSA subpool 227 with KEY=6. Pageable buffer pools are allocated in extended CSA subpool 231 with KEY=6.

Table 60. VTAM buffer pools			
Buffer pool	Use	Default storage type	Recommendations and requirements
APBUF	Used to provide fixed common storage if VTAM fails to obtain other storage from the operating system. For buffers not related to I/O.	Fixed	IBM-supplied values are appropriate for most systems.
BSBUF	Used to maintain session information for peripheral nodes for which VTAM performs boundary function. One buffer is required for each boundary LU session (SSCP-PU, SSCP-LU, LU-LU).	Fixed	Set base number to the average number of concurrent boundary LU sessions.
CRA4	Used for scheduling and error recovery.	Pageable	IBM-supplied values are appropriate for most systems.
CRA8	Used for scheduling and error recovery.	Pageable	IBM-supplied values are appropriate for most systems.
CRPLBUF	RPL-copy pool. In general, one buffer is required for each VTAM application program request until the operation is complete.	Pageable	Monitor output from the DISPLAY BFRUSE command to ensure that the number of buffers specified is adequate for your system.

Table 60. VTAM buffer pools (continued)

Buffer pool	Use	Default storage type	Recommendations and requirements
IOBUF	Used for input/output data. Every PIU that enters or leaves VTAM resides in an I/O buffer. This pool is 31-bit addressable.	Fixed	<p>Monitor output from the DISPLAY BFRUSE command to ensure that the number of buffers specified is adequate for your system.</p> <p>Ensure that the largest MAXBFRU value is less than <i>xpanpt</i> minus <i>slowpt</i>. Also ensure that <i>bufsize</i> is greater than or equal to the UNITSZ operand on the HOST definition statement in the NCP.</p> <p>Note: See the z/OS Communications Server: SNA Resource Definition Reference for additional requirements and restrictions concerning the IOBUF pool.</p>
LFBUF	<p>One buffer is required for each active application program with an EAS value (on APPL definition statement) less than 30. If the EAS value is greater than 30, this information is contained in SFBUF.</p> <p>One buffer is required for each TSO user who is logged on.</p>	Fixed	IBM-supplied values are appropriate for most systems.
LPBUF	Used for scheduling and audit trail (error recovery). One buffer is required for each active VTAM process.	Pageable	Set <i>baseno</i> to 9, and set <i>xpanno</i> in the range 4–6 to force 2-page expansion.
SFBUF	Used to contain application program information and LU blocks. One buffer is required for each active application program with an EAS value (on APPL definition statement) greater than or equal to 30.	Fixed	IBM-supplied values are appropriate for most systems.
SPBUF	Used for large message (LMPEO) requests. One buffer is required per concurrent LMPEO send request.	Pageable	IBM-supplied values are appropriate for most systems.

Table 60. VTAM buffer pools (continued)			
Buffer pool	Use	Default storage type	Recommendations and requirements
TIBUF	Used to perform input/output operations for CSM-capable protocols. This pool is in 31-bit storage.	Fixed	<p>IBM-supplied values are appropriate for most systems. The default value for this pool is set at a conservative value in case the functions that use this pool are not used. If using the functions that use this pool, use the DISPLAY net,BFRUSE command to monitor usage and then set the BASENO for the pool at the normal high period usage.</p> <p>For HPR and IP - Defining more initial buffers and larger buffer extents may help increase throughput. Defining too few may cause an IP retransmit.</p>
T1BUF	Similar to the TIBUF but larger. Used as a packing buffer by HiperSockets accelerator and QDIO. Also used to contain the HPR headers and the media, IP, and UDP headers for an Enterprise Extender connection.	Fixed	<p>IBM-supplied values are appropriate for most systems. The default value for this pool is set at a conservative value in case the functions that use this pool are not used. If using the functions that use this pool, use the DISPLAY net,BFRUSE command to monitor usage and then set the BASENO for the pool at the normal high period usage.</p> <p>For HPR/IP - Defining more initial buffers and larger buffer extents might help increase throughput. Defining too few might cause an IP retransmit.</p>
T2BUF	Similar to the T1BUF but larger. T2BUFs are used exclusively for HiperSockets accelerator and QDIO when the T1BUF is perceived to be of insufficient size to pack all the headers and data. Also used by HPR for all retransmissions over HiperSockets and QDIO.	Fixed	<p>IBM-supplied values are appropriate for most systems. The default value for this pool is set at a conservative value in case the functions that use this pool are not used. If using the functions that use this pool, use the DISPLAY net,BFRUSE command to monitor usage and then set the BASENO for the pool at the normal high period usage.</p> <p>For Enterprise Extender connections using QDIO/iQDIO device drivers, defining more initial buffers and larger buffer extents might help increase throughput. Defining too few might cause an IP retransmit.</p>

Table 60. VTAM buffer pools (continued)			
Buffer pool	Use	Default storage type	Recommendations and requirements
XDBUF	Used for physical I/O to VTAM peripheral nodes for connection activation (for example, XID exchange processing). This pool is 31-bit addressable.	Fixed	IBM-supplied values are appropriate for most systems. Because this buffer pool handles the concurrent activation of physical units, you should specify a small size and enable dynamic buffering. The buffer pool can then expand to accommodate the concurrent activation and can shrink after the resources are active.

Buffer pool allocation

VTAM provides two types of buffer pool storage allocations, basic and dynamic.

Basic allocation

The amount of space that you reserve for the buffer pool when VTAM is started.

Dynamic allocation

The process by which VTAM temporarily increases the size of a buffer pool when there are heavy demands for space in that pool.

While basic allocation sets the base size of your buffer pools, dynamic expansion enables a buffer pool to be expanded temporarily during periods of heavy demand. Its use can greatly increase the efficiency with which VTAM uses storage, particularly for I/O buffers.

Without dynamic expansion of a pool, basic allocation parameters must be specified large enough to meet the greatest possible demands on the pool. With dynamic expansion, smaller basic allocation values can be specified, and peak demands on the pool can be met with dynamic expansion. Dynamic expansion is strongly recommended for most buffer pools because the peak demand can vary considerably from the normal demand.

Setting buffer pool allocations

To set buffer pool space, use the buffer pool start options:

```
poolname=(baseno,bufsize,slowpt,F,xpanno,xpanpt,xpanlim)
```

For a description of these operands, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

The first four operands of the start option (*baseno*, *bufsize*, *slowpt*, and *F*) establish the base size and the characteristics of the buffer pool. The last three operands (*xpanno*, *xpanpt*, and *xpanlim*) specify how dynamic expansion of the buffer pools is to be performed.

Set basic allocation for each buffer pool when VTAM is started. If you do not specify a start option parameter for a buffer pool, VTAM uses an IBM-supplied value for the missing parameter. The IBM-supplied values are not necessarily the most efficient values for your system, and they are not necessarily compatible with any dynamic allocation specifications you might make.

How dynamic expansion operates

1. You enable dynamic expansion by coding the *xpanno* and *xpanpt* operands on the buffer pool start option.
2. The buffer pool size reaches the point specified by *xpanpt*.

3. VTAM acquires more buffers in blocks to accommodate the increased demand. The size of the blocks is specified by *xpanno*. VTAM acquires buffers as needed until the pool size reaches *xpanlim* (if specified).
4. Contraction is determined by *xpanno* and *xpanpt*. The *xpanno* value used is the *xpanno* value that was specified when VTAM was started, but rounded upward to the number of buffers that fills the nearest whole page of storage. If the buffer pool is the IO00 or TI00, when the number of available buffers is greater than or equal to $(3 \times xpanno) + xpanpt$, VTAM verifies that the buffers that it acquired in previous expansions of the pool are not in use. For all other pools, the value is $(2 \times xpanno) + xpanpt$. (If they are not, VTAM releases these buffers in blocks (as when they are acquired). If any of the buffers in a block are in use, VTAM does not release that block of buffers.

Figure 142 on page 540 shows the structure of a pool after basic allocation A and after one dynamic expansion of the pool B.

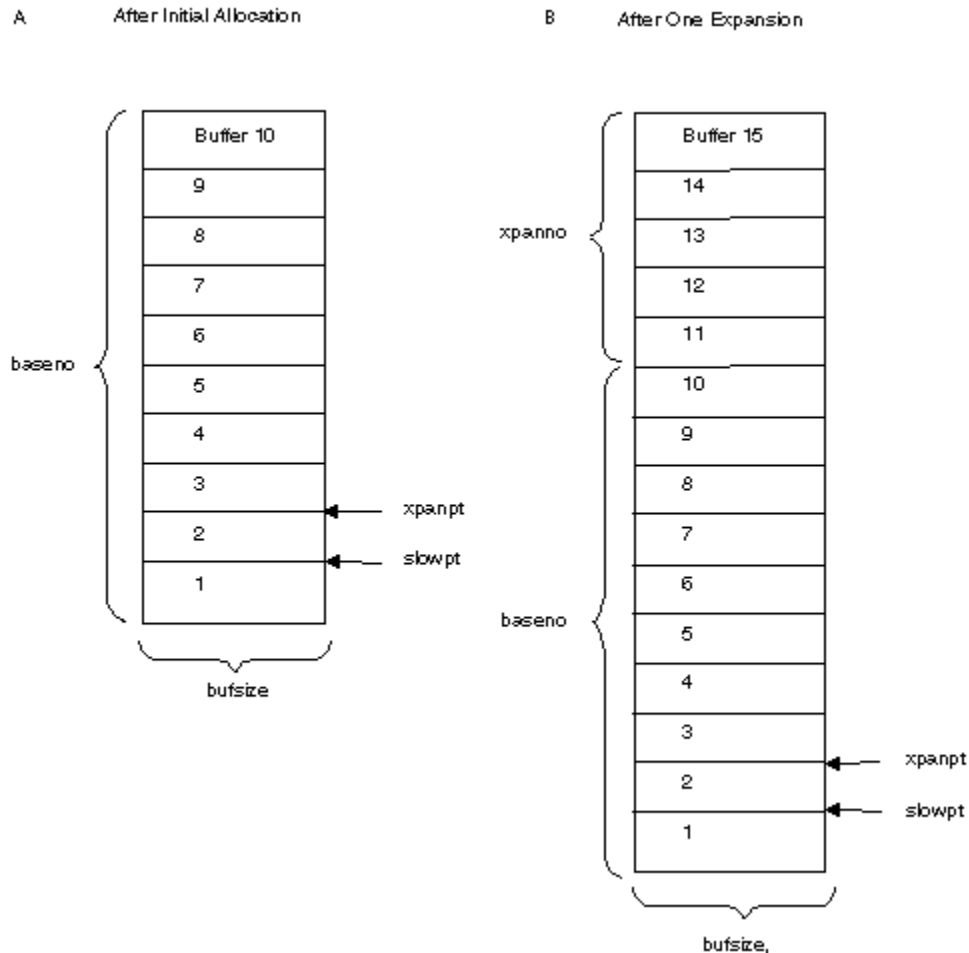


Figure 142. Buffer pool after initial allocation and after one expansion

A

This example shows a buffer pool for which the start options were specified as `poolname=(10,bufsize,1,,5,2,xpanlim)`. After initial allocation, the pool contains ten buffers (*baseno*=10), the length in bytes of each buffer is *bufsize*, the slowdown point (*slowpt*) is one, the expansion size (*xpanno*) is 5 (assume that five buffers fill one page of storage), and the expansion point (*xpanpt*) is 2. The maximum allowed size of this buffer pool is determined by the value *xpanlim*.

B

After one expansion, there are 15 buffers in the pool. Each of the five additional buffers has a length of *bufsize* and the same expansion point and slowdown point as before.

Guidelines for dynamic expansion

The following guidelines help you determine appropriate buffer pool values for your system.

If *xpanno* or *xpanpt* has a value of 0, no expansion occurs.

- Set *baseno* for the I/O buffer to the steady state value plus *xpanpt*.

You can determine steady state by issuing the DISPLAY BFRUSE operator command. CURR TOTAL minus CURR AVAIL gives the number of buffers in use. Repeat this procedure several times, and take the average to determine the steady state value. Be sure to include in your calculation displays taken during times of peak network usage.

Any additional buffers needed can be obtained using dynamic expansion.

- Define dynamic expansion so that pools are expanded one page at a time.

To do this, define *xpanno* to be the number of buffers that fit on one page. To calculate *xpanno* for I/O buffers, use [Table 62 on page 543](#), and do the following steps:

1. Find the *bufsize* range in the table that includes the *bufsize* value that was specified for your system.
 2. Round up the current *bufsize* to the largest even number that appears in that range. (Some channel-attached SNA devices require that the I/O *bufsize* be an even number of bytes. Therefore, VTAM rounds up *bufsize* to an even number if an odd number is specified.)
 3. For one-page expansion, specify an *xpanno* value that is equal to the number in the left column of this row (the number of I/O buffers for each page).
- Tune the expansion increments (*xpanno*) and expansion points (*xpanpt*) to keep CPU overhead from dynamic buffering low.
 - Expansion and contraction of buffer pools increase the amount of real storage that VTAM requires. If the buffer pools are continually expanding and contracting, you should change the values specified for *xpanno* and *xpanpt*. You can determine what values to specify by monitoring the output of the DISPLAY BFRUSE operator command or the SMS buffer trace and by modifying the values for *xpanno* and *xpanpt* accordingly.
 - If the expansion point is set too small, VTAM might not be able to expand the buffer pool fast enough if there is a rapid demand for buffers. Setting the expansion point too small also causes the buffer pool to contract frequently.

Message IST561I might indicate that storage is temporarily unavailable because of this rapid demand. Adjust the *xpanpt* and *xpanno* values to eliminate this problem.

The expansion point for I/O buffers should be larger than the largest single request for buffers. Following are examples of the largest single request:

- MAXBFRU allocation
- NetView session monitor trace buffer size
- JES/NJE TPBFSIZ
- If excessive expansion and contraction are a problem (especially for LPBUF and CRPLBUF), define these buffers to expand two pages at a time.

The buffer pool can expand and contract too frequently if the expansion increment and expansion point are not large enough to handle normal fluctuation in the pool. The I/O buffer pool is prone to this type of thrashing from application programs that send very large messages. VTAM always rounds up an expansion increment to fill a multiple of a full page.

- Initially allocate one page to SFBUF.

Any elements (where an element consists of a buffer, whose size is rounded to a multiple of 8 bytes, plus a 16-byte header) that are required over a page should be controlled by dynamic expansion.

- Set CRPLBUF.

In environments where many sessions come up simultaneously because of either the VARY LOGON operator command or the LOGAPPL operand on a definition statement, you might need to preallocate enough CRPLBUF buffers to handle the influx of requests.

- Set LFBUF.

LFBUF is used only for EAS values (on APPL definition statements) less than 30. Because this pool is not used often, a *baseno* value of 2 is recommended. If you are using TSO, a buffer from the LFBUF buffer pool is used for each TSO user who is logged on. The pool should be expanded one page at a time.

- Specify *xpanlim* for the IOBUF pool.

xpanlim is useful in constraining the I/O buffer pool from excessive growth, particularly when a logical unit attempts to flood VTAM with requests (HOT I/O). Typically, the logical unit is malfunctioning; however, the logical unit may be functioning correctly but may not have had proper session parameters specified. Specifying *xpanlim* allows VTAM to detect such sessions before complete CSA exhaustion and, if the HOTIOTRM Start Option is also specified, terminate such sessions without the system being impacted.

To set *xpanlim*, issue DISPLAY BFRUSE,BUFFER=IO find the peak usage, and set the limit somewhat above that peak. Because *xpanlim* is specified in units of 1024 (K) bytes, you will need to convert buffers to KB. Use Table 62 on page 543 to determine the number of I/O buffers of a given size that fit on a 4K page. Use the equation $K = (\text{number of buffers} / \text{number of buffers per page}) * 4$.

Tip: You can modify the *xpanlim* value after VTAM startup using the MODIFY BFRUSE command. See [z/OS Communications Server: SNA Operation](#) for information about this command.

Also, bear in mind that in order for the expansion limit to be meaningful, CSA must be available were the expansion limit to be reached and VTAM must not reach any CSA limit value that has been specified. In other words, ensure that you have adequate CSA available in the system and your CSA limits allow VTAM to actually allocate buffers up to your expansion limit.

Note: In Table 62 on page 543, where it lists a range of buffer sizes for a given buffers per page value, whenever possible use the highest value in the range. Some devices have specific I/O buffer size requirements that may prevent using the highest value.

Table 61 on page 542 lists the number of buffers per page for the different buffer pools.

Table 61. Number of buffers per page		
Buffer pool	Default storage type	Buffers per page
APBUF	Fixed	56
BSBUF	Fixed	14
CRA4	Pageable	1
CRA8	Pageable	NA
CRPLBUF	Pageable	25
LFBUF	Fixed	30
LPBUF	Pageable	2
SFBUF	Fixed	32
SPBUF	Pageable	21
TIBUF	Fixed	5
T1BUF	Fixed	3
T2BUF	Fixed	2
XDBUF	Fixed	5

Table 62. I/O buffer size and number of buffers per page

Number of buffers per page	Range of buffer sizes for MVS (with 4K paging)	Range of buffer sizes for MVS using data encryption facility (with 4K paging)
1	1958 - 3992	1951 - 3985
2	1270 - 1957	1263 - 1950
3	934 - 1269	927 - 1262
4	726 - 933	719 - 926
5	590 - 725	583 - 718
6	494 - 589	487 - 582
7	422 - 493	415 - 486
8	358 - 421	351 - 414
9	318 - 357	311 - 350
10	278 - 317	271 - 310
11	256 - 277	256 - 270
Note: The buffer size in the table is the size coded on the IOBUF start parameter. For information about the calculations used in arriving at these values, see the z/OS Communications Server: SNA Resource Definition Reference .		

Dynamic expansion and storage

As buffer pools expand, expansion uses virtual storage that can potentially impact an operating system that has virtual storage constraints. VTAM use of virtual storage can be limited by the CSA start options and the virtual storage allocated to VTAM address space.

All buffer pools are defined in virtual storage above the 16 MB address line in extended virtual storage. Therefore, the size of the buffer pools does not usually impact virtual storage constraints.

You can use the SONLIM start option to limit the amount of fixed I/O buffer storage available for session outage notification messages. This controls VTAM use of fixed (real) storage. The default (60 percent) should be sufficient for most environments. This percentage is calculated from the base number of buffers specified when VTAM is started but does not include buffers from dynamic buffering in its calculations. Therefore, the default value of 60 percent is normally sufficient if you do not preallocate most of your I/O buffers at start time and if you use dynamic buffering. If you preallocate most of your I/O buffers at start time, you should consider overriding the VTAM default with a smaller percentage.

I/O buffers and application program data transfer

Data that is transferred between a VTAM application program and another logical unit in the network passes through VTAM I/O buffers. VTAM temporarily holds the data until an application program requests it or until it can be sent to a logical unit.

[Figure 143 on page 544](#) shows how VTAM uses buffers for input operations.




Figure 143. How VTAM uses input buffers

1. A logical unit in the network can send data to an application program independent of the application program requesting the data. The data is placed in a buffer in the fixed buffer pool of data buffers.
2. One of the following occurs:
 - a. If the application program has an outstanding request for data from the logical unit, VTAM moves the data from the buffer in the fixed buffer pool to the specified application program data area. Later, data can be paged out of main storage by the operating system.
 - b. If the application program does not request data from the logical unit, VTAM places the data in a data space.
3. When the application program issues an input request, the operating system pages the data into main storage (if necessary) and VTAM moves the data to the specified application program data area.

In general, an application program should be written so that at least one request for input is always outstanding. This reduces the chance of VTAM using excessive amounts of an application program's dataspace storage.

Figure 144 on page 544 shows how VTAM uses buffers for output operations.

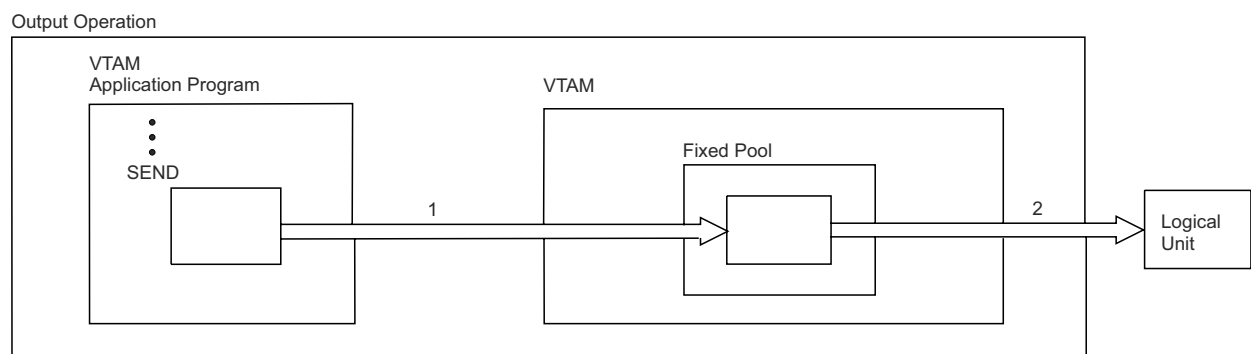


Figure 144. How VTAM uses output buffers

1. The application program requests data transmission. When VTAM has sufficient buffers in the fixed pool, it moves the data to the fixed pool.
2. VTAM sends the data to the NCP or the logical unit, freeing the buffer.
3. If HPR is used, VTAM holds onto the I/O buffers in pageable storage until the other RTP endpoint acknowledges receipt of the data.

If you notice that your I/O buffers have expanded and you are running out of CSA, it is possible that an application program is not pacing data. Investigate the application programs that have a lot of traffic moving through the I/O buffers to determine whether pacing is being performed. For general information about pacing, see [“Session-level pacing”](#) on page 217. For specific information about application program pacing, see [“Application program pacing”](#) on page 223.

HOT I/O detection/termination

VTAM has functions to both detect and terminate sessions deemed to be using excessive amounts of the I/O buffer pool. The I/O buffer pool is the storage pool that contains session data. Typically, these sessions are undergoing some sort of malfunction causing an unrestrained transmission of data into VTAM. A failure of VTAM to detect and terminate these out-of-control sessions can result in a VTAM or system outage because of CSA exhaustion. These sessions are commonly referred to as HOT or HOT I/O.

HOT I/O session detection

VTAM scans the entire I/O buffer pool for sessions using excessive I/O buffers under two conditions:

- When a successful expansion of the I/O buffer pool brings the total size of the pool to 80 percent (or more) of the expansion limit specified (in the IOBUF statement). Note that this happens for each expansion in which the size of the pool is greater than the above mentioned threshold.
- When an expansion of the I/O buffer pool is unsuccessful (for any reason). In this case message IST154I is issued containing the expansion failure reason code. Note that if an expansion limit is not specified, then this is the only condition in which a HOT I/O scan is performed.

When one of the above conditions occurs, VTAM scans the entire I/O buffer pool looking for sessions using greater than 10 percent of the current size of the I/O buffer pool. Any session found to be using greater than 10 percent of the pool will have an IST930I message issued for it. This message contains the session partner names and the percent of the pool in use by the session.

Specifying an expansion limit for the I/O pool is advantageous because VTAM can detect and correct HOT I/O conditions before a complete exhaustion of common service area (CSA). There are disadvantages involved when an incorrect value has been specified. Consequently, it is recommended that the following guidelines be carefully evaluated so that a proper value may be specified.

If the expansion limit has been specified too low, one or more of the following condition may occur:

- VTAM uses extra CPU cycles doing HOT I/O scans. This primarily occurs when the expansion limit is specified so low that under normal workloads the size of the I/O buffer pool is already greater than 80 percent of the expansion limit. The amount of CPU cycles that VTAM expends doing a HOT I/O scan is dependent on the size of the I/O buffer pool.
- VTAM cannot expand the I/O buffer pool. This occurs when the expansion limit has been specified so low that VTAM cannot expand the I/O pool to accommodate a normal workload increase because of it reaching the expansion limit. VTAM may be adversely affected by a failure to obtain more I/O buffers.

If the expansion limit has been specified too high, one or more of the following condition may occur:

- VTAM does not detect a HOT I/O condition until after CSA has been exhausted or limited by some other factor (such as the CSALIMIT start option). The expansion limit was specified so high that the size of the pool never reached the 80 percent threshold and consequently VTAM never scanned for HOT I/O until an expansion failed. This works as if there is no expansion limit specified. Consequently, VTAM and the system may be adversely affected by this condition.
- A VTAM or system outage occurs because of an exhaustion of CSA. This primarily occurs when VTAM CSA use is not also limited with the CSALIMIT parameter. If VTAM is not limited in the amount of CSA

it can acquire and a HOT I/O situation occurs, VTAM may use so much of CSA that the system (GCS or MVS) is adversely affected.

Tip: If the expansion limit is found to be too high or too low, you can modify the *xpanlim* value using the MODIFY BFRUSE command. See [z/OS Communications Server: SNA Operation](#).

HOT I/O session termination

Using the HOTIOTRM start option, you can set the threshold at which VTAM will terminate "hot" (out-of-control) sessions. HOTIOTRM specifies the percentage (10–99) of the current size of the I/O buffer pool that a single session must have allocated to it to cause VTAM to take corrective action. If HOTIOTRM is allowed to default to 0, VTAM will not terminate sessions based on I/O buffer pool usage.

If HOTIOTRM is specified with a value in the range 10–99, and 1 or more sessions have been identified by VTAM as using a percentage of the I/O buffer pool greater than or equal to that value, VTAM will then attempt to stop the buffer pool expansion with one of the following actions:

- Deactivate the LU when the session type is SSCP-LU.
- Deactivate the PU when the session type is SSCP-PU.
- Terminate all sessions between the two logical units when the session type is LU-LU.

The above action will be taken if all of the following criteria are met:

- The percentage of the pool in use by the session is greater than or equal to the HOT I/O session termination threshold as specified by the HOTIOTRM start option.
- The subarea of one or both session partners must be the subarea in which the HOT I/O is being detected.
- When the session type is SSCP-PU, the PU is not an NCP.

If all of the above conditions are true and it is an LU-LU session, VTAM terminates all sessions between the two session partners. It then issues message IST1099I, indicating that the sessions have been terminated.

If all of the above conditions are true and it is a SSCP-LU or SSCP-PU session, VTAM deactivates the LU or PU. It then issues message IST1098I, indicating that the LU or PU has been deactivated.

Maximizing coattailing

To maximize coattailing, you need to transfer more data either inbound to the host or outbound without generating an attention interrupt. If more messages are transferred in or out of the host than the number of READ or WRITE channel programs issued, coattailing is taking place. The effect is to transfer more data in a single channel I/O operation. Coattailing can provide greater throughput between a host and an SNA controller; however, response time can increase.

The amount of coattailing that you can achieve is directly related to the message traffic in your network. However, you can influence the amount of data that is available for a data transfer operation.

Coattailing allows more PIUs to be queued for transmission before a channel program is issued. This entails a trade-off between response time and host/controller cycles. Specifying a large value for coattailing delay usually results in a larger number of PIUs being accumulated before a channel program is issued. However, PIUs will experience a longer wait time, and consequently, the overall response time might increase. The benefits of specifying a coattailing delay are dependent on the volume, distribution, and predominant direction (if any) of traffic flow.

The coattailing delay is implemented based on PIU arrivals. The arrival of a PIU causes VTAM to compute the elapsed time since the beginning of the delay interval and determine if the coattailing delay has been exceeded. This has implications for the benefits of coattailing:

- At low traffic volumes, PIU arrivals might be infrequent. This could cause a PIU to stay in the queue until the next PIU arrives or the backup three-second timer expires. Setting a nonzero coattailing delay could delay transmission of a PIU up to three seconds. The coattailing delay should be set to 0 for low traffic volumes.

- At high traffic volumes, sufficient numbers of PIUs are queued while a previous channel program is running. This results in "automatic" coattailing. Specifying a coattailing value greater than 0 might not result in greater throughput or reduced host processing cycles.
- Coattailing can be beneficial at high processor utilization with moderate traffic volumes. In this case, tuning the coattailing delay parameters might result in reducing the impact on host processor cycles and provide improved throughput.
- For predominantly unidirectional data traffic, specifying a coattailing delay for traffic in the nonpredominant direction might result in larger overall delays and reduced throughput. For example, file transfers with application level acknowledgments for segments of files transferred might see lower throughputs with a nonzero coattailing value because the acknowledgment PIUs can be delayed for up to three seconds if no other PIU triggers the coattailing timer.

In general, transaction processing applications using 3270 data streams are characterized by bidirectional traffic with relatively small messages. In this case, a nonzero value of coattailing delay might be beneficial at some traffic volumes. Newer client-server applications are usually characterized by fewer and larger messages. These applications will see minimal benefit with nonzero coattailing delays.

Controlling outbound coattailing

By controlling the channel delay timer for VTAM and the attention interrupt delay timer, you can affect outbound coattailing in your system.

VTAM saves PIUs for transfer over a channel until one of the following events occurs:

- Data flowing at transmission priority two (TP2) is available.
- A virtual route pacing response must be transmitted.
- An attention interrupt from an NCP or an SNA cluster controller (for example, the 3174) occurs.
- The PIU queue for VTAM reaches its limit.

For each VTAM channel delay interval, VTAM calculates the queue depth (QDPTH tuning statistic) by multiplying the number of PIUs transferred during the channel delay interval by a percentage (currently 75 percent).

- The channel delay timer or the backup three-second timer expires.

Controlling inbound coattailing

To affect inbound coattailing, delay the VTAM WRITE operation. This is effective because VTAM normally reads data after transferring data to the SNA controller or to another VTAM. By delaying VTAM longer before it performs a WRITE operation, more data can potentially accumulate in the SNA controller or other VTAM host, which can be read in the same I/O operation as the WRITE channel program. You can control the time interval during which VTAM buffers low priority data before transferring it over the channel to another VTAM, NCP, or SNA cluster controller by coding the DELAY operand on the appropriate VTAM definition statement.

DELAY operand

When defining a channel, consider the effects of the DELAY operand on the LINE definition statement. This operand specifies the maximum amount of time that VTAM queues an outbound PIU with transmission priority 0 or 1 before transmitting it over the channel. When the PIU is eventually sent, it and the queued PIUs following it are sent as a block.

When you specify DELAY=0, VTAM sends each PIU over the channel as it arrives on the queue. When a nonzero value for DELAY is specified, the blocking of low-priority PIUs can save host processor instructions. However, a high channel delay time can cause increased response time because some of the PIUs are delayed before being sent.

Occasional bursts affect throughput only for the next second or two. However, a continuing series of bursts can severely affect capacity if there is little nonchannel-to-channel activity. A few sessions of

interactive traffic are more likely to accentuate this condition than a large number, when traffic on the channel is likely to be more randomly distributed.

You might want to increase the DELAY operand value over that of the default to:

- Decrease host processor overhead when response time is less important.
- Take advantage of a smoothing effect on the queuing that can occur with larger values, an effect that can be useful for application programs with sessions conducted in bursts.

You can specify different DELAY values for different channel-attached devices. The following are general guidelines for picking an appropriate DELAY value for a particular channel-attached device.

If you want to experiment with DELAY values greater than the default, you can activate in succession a series of major nodes with increasing values beyond the default.

MAXBFRU operand

For channel-to-channel connections, the MAXBFRU operand defines a constant sized buffer in increments of 4K pages. The value does not fluctuate throughout the life of the connection. This buffer must be defined as large as the largest PIU to cross the channel-to-channel connection. Also, if a nonzero DELAY value is chosen and you want to transfer multiple RUs in a single I/O operation, the MAXBFRU value must also be large enough to handle multiple RUs.

After the data transfer operation is complete, the PIUs are deblocked into VTAM buffers as defined in the IOBUF start option. The VTAM buffers are allocated in 4K increments, so you might have to allocate more I/O buffers. If you are using dynamic buffer expansion, you should monitor the I/O buffers pool to ensure that there is not excessive buffer pool expansion and contraction because of these channel operations.

The maximum PIU size across a connection is defined by MAXBFRU times IOBUF size. For channel-attached cluster controllers, the maximum PIU size is defined by MAXBFRU times (IOBUF size minus 16).

If all LU-LU sessions end in a physical unit that provides segment assembly, path information unit (PIU) segmentation occurs across the link. If any LU-LU sessions end in a physical unit that does not provide segment assembly, the maximum PIU size must be large enough to handle RU sizes for these sessions. If the LU-LU session cannot have segments, either increase the value of the MAXBFRU operand or the I/O buffer sizes, or adjust the RU size in the MODEENT macroinstruction for the logical unit.

Coattailing for SNA controllers

The amount of coattailing of messages that occurs between a host and a channel-attached NCP or SNA controller is affected by:

- I/O buffer size
- The number of buffers allocated to receive incoming messages at the host
- The amount of time that VTAM and the channel-attached device allow to elapse before sending buffered messages

For NCP Version 4 Release 2 and later releases, the UNITSZ value specified in the NCP HOST definition statement is no longer used. If coded, the value is ignored. The value of VTAM IOBUF size is determined by NCP during XID exchange. For NCP releases before 4.2, the value of UNITSZ can affect VTAM, NCP, and network performance. This section provides guidelines for specifying this value.

The initial amount of storage you allocate for the I/O buffer, and the *xpanno* value you specify, depends on the *bufsize* specification. This, in turn, depends on the UNITSZ value specified in both the NCP HOST definition statement and the buffer pool start option for the I/O buffer. In the case of channel-attached cluster controllers, *bufsize* depends only on the buffer pool start option for the I/O buffer. The value in the VTAM start option is used by both VTAM and the NCP when they exchange information. The *bufsize* value entered as part of the buffer pool start option for the I/O buffer is interpreted by VTAM as the UNITSZ value used, except for channel-attached cluster controllers, where UNITSZ is IOBUF minus 16. The I/O buffer size, as it appears on the console after a DISPLAY BFRUSE operator command, is greater than the *bufsize* that was specified in the buffer pool start option because VTAM adds control information to each I/O buffer.


Guidelines for setting DELAY

To specify the amount of time that elapses before VTAM or NCP sends buffered data, use the DELAY operand on NCP and VTAM definition statements. The DELAY operand on the NCP BUILD definition statement controls how long the NCP buffers data; it is the elapsed time between the receipt of the first inbound message and the presentation of an attention interrupt to the host. This capability can affect your strategy for increasing coattailing use of the channel.

The DELAY operand that controls how long VTAM buffers data can be specified on the following statements:


- PCCU definition statement for NCPs not defined in a channel-attachment major node
- GROUP, LINE, or PU definition statements for NCPs defined in a channel-attachment major node
- PU definition statement for local (channel-attached) SNA controllers

The following three examples illustrate the use of the DELAY operand.



In this example, four messages have come in over a period of time. Because DELAY is 0, each message has resulted in an attention interrupt being presented to the host. In all, four attention interrupts have been presented, and no coattailing has taken place.

Figure 145. Effect of DELAY time on coattailing - example 1



In this example, DELAY is equal to t . This time period is large enough to allow four messages to arrive at the NCP. The time t has expired, and the NCP presents an attention interrupt. VTAM reads all four messages, having received only one attention interrupt.

Figure 146. Effect of DELAY time on coattailing - example 2




Figure 147. Effect of DELAY time on coattailing - example 3

For NCP, the DELAY operand allows time for more than one message to arrive in the NCP before an attention is presented to the host.

The following events cause PIUs on the queue to be sent before the DELAY time elapses:

- VTAM receives an attention interrupt from NCP.
- VTAM or NCP reaches the MAXBFRU limit.
- VTAM sends virtual route pacing responses.
- VTAM sends transmission priority two traffic.
- VTAM sends APPN network priority or highest priority traffic (to a local SNA PU).

The default value for the DELAY operand is 0 (for example, no coattailing). Take the default under the following conditions:

- You have a low rate of transactions that use the channel connection, and response time for interactive transactions is important.
- You have a high rate of transactions that use the channel connection, and response time for interactive transactions is important.
- You have predominantly unidirectional traffic (file transfer) with little other activity in the opposite direction.
- You are benchmarking single or few sessions of bulk data traffic (file transfers, backups, and so on).
- You have channel traffic with frequent bursts of SENDs (such as from an application program that makes inquiries of the other host in rapid succession).

To analyze inbound coattailing for channel-attached SNA communication or cluster controllers, compare the number of times that NCP signals VTAM that it has data to send (ATTN tuning statistic) to the number of times that VTAM issues a channel READ program to read data (CHRD or RCH tuning statistic). A less accurate method is to divide the number of inbound PIUs transferred by NCP (IPIU tuning statistic) by the number of read channel programs issued by VTAM (CHRD or RCH). This calculation gives you the average number of PIUs read from the channel-attached device with each channel program, which is an indicator of coattailing.

Note: For a 3274 to clear its buffers after a data transfer operation, the 3274 signals to VTAM that it has more data to send even though there is no data. This causes the RDATN tuning statistic to equal the CHRD or RCH tuning statistic, and because RDATN is included in the ATTN tuning statistic, the ATTN value is normally double the CHRD or RCH value.

To analyze outbound coattailing for channel-attached SNA communication or cluster controllers, divide the number of outbound PIUs (OPIU tuning statistic) by the number of WRITE channel programs VTAM issues (CHWR or WCH tuning statistic). This calculation gives you the average number of PIUs transferred by VTAM with each channel write program, which is a rough estimate of the effectiveness of coattailing.

Specifying a nonzero attention DELAY value to activate attention DELAY does not automatically ensure coattailing. To ensure coattailing, all of the following must be true:

- The installation must have a transaction rate high enough to allow more than one message to accumulate in the DELAY time period in the NCP or the host.
- The MAXBFRU value must be large enough to allow more than one message to be sent to the host.
- The host processing speed has an effect on coattailing; the faster host processors reduce coattailing slightly. An attention interrupt occurs when MAXBFRU is reached or when the DELAY time period expires.

Begin by setting DELAY equal to or greater than 0.2. A 0.2-second delay has little effect on response time, but if the traffic speed is approximately 15 transactions per second or greater, coattailing occurs.

Setting the MAXBFRU operand

Using the MAXBUF operand, you can specify the number of VTAM I/O buffers allocated for inbound data transfer. You need to set MAXBFRU on both the NCP HOST definition statement and the VTAM GROUP or LINE definition statements for channel-attached NCPs or SNA controllers. When VTAM creates a read channel program, it must always have MAXBFRU number of I/O buffers available for reading from the channel-attached device. VTAM always attempts to hold in reserve MAXBFRU number of I/O buffers for the next read from the channel-attached device.

VTAM and NCP do not support segmenting PIUs across channel connections between them. Therefore, ensure MAXBFRU times IOBUF size accommodates the largest PIU that NCP can send to VTAM. NCP informs VTAM the maximum size PIU that VTAM can send in its XID2 (this is based on TRANSFR*BFRS-18). VTAM also honors the MAXDATA keyword coded on the PCCU definition and will send PIUs to NCP no larger than the smaller of the MAXDATA keyword or what VTAM received from NCP in the XID2. Therefore, you should ensure that MAXDATA and TRANSFER*BFRS-18 is larger than the largest PIU VTAM can send to NCP.

In addition, choose MAXBFRU so that more messages can be coattailed. The MAXBFRU value determines when an attention interrupt is signaled to the host. MAXBFRU=5 is a good starting point. There should be enough buffers available to transfer all the messages likely to arrive in the DELAY time period. (The use of MAXBFRU=5 assumes one inbound message fits in one I/O buffer. More details on I/O buffer size are given in the following paragraphs.) You can use tuning statistics to verify that MAXBFRU is sufficient for coattailing. Using the example in Figure 150 on page 553, for instance, MAXBFRU could have been set to 2. Instead, it was set to 6 to maximize coattailing.

MAXBUFU formulas

The number of buffers required for a message (RU) is determined by the following formula:

$$(29 + \text{RU size}) \div \text{bufsize}$$

Note: This formula applies to FID4 (subarea connections).

The number 29 is obtained by adding the size of the transmission header (26) to the size of the request/response header (3).

For channel-attached cluster controllers, the number of buffers required for a message is determined by the following formula:

$$(13 + \text{RU size}) \div (\text{bufsize} - 16)$$

The number 13 is the sum of the channel link header (4) plus the transmission header (6) plus the request/response header (3).

The *bufsize* is the buffer pool value you use to start VTAM. If you are unsure about this value, look at your start options for the IOBUF buffer size specification.

The MAXBFRU value you use must be greater (for NCPs and PU2.0 channel attached devices) than or equal to the number of buffers required for the largest message (RU) your system receives. For T2.1

channel-attached devices, the value should be greater, but does not have to be. This minimizes the number of attention interrupts needed to process large messages. Use the preceding formula to calculate the lower limit for MAXBFRU, given the size of the largest message.

Guidelines for setting I/O buffer size

Choosing the ideal I/O buffer size helps minimize host storage requirements and host processor usage.


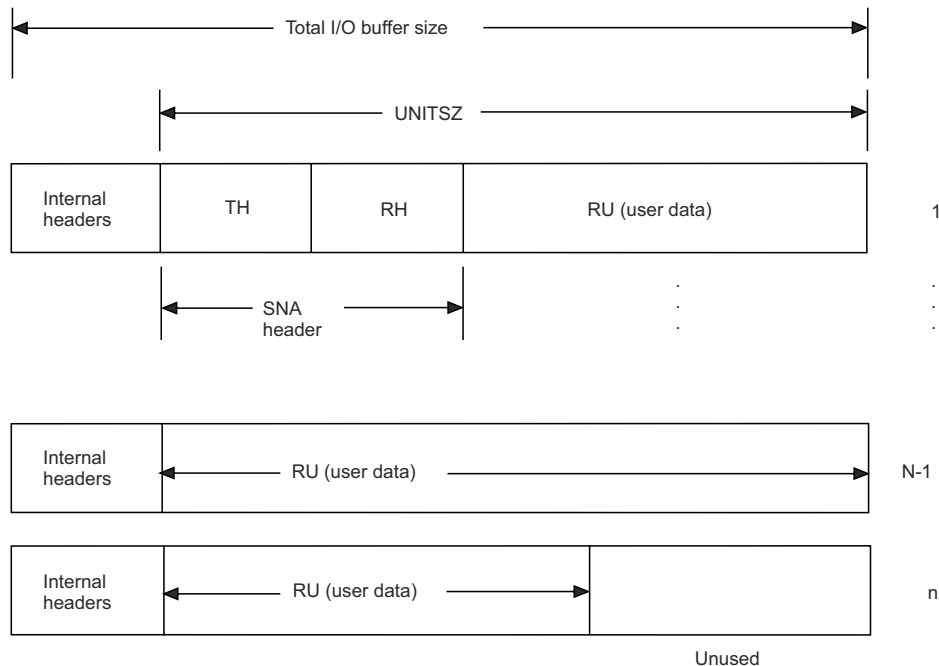


Figure 148. General I/O buffer format

Figure 148 on page 552 illustrates UNITSZ. UNITSZ is not the size of an I/O buffer; it is a portion of the buffer size. It includes an SNA header field. In the remaining discussion, the I/O buffer size is referred to as the UNITSZ.

Note:

One buffer size is used for both inbound and outbound data transfers.



Note: Multiple I/O buffers are used where the length of outbound data is greater than UNITSZ minus the SNA header.

Note: A similar concept applies to inbound data.

Figure 149. Using multiple I/O buffers to transfer single message

Figure 149 on page 552 illustrates how a message that is larger than UNITSZ spans more than one buffer. The message can be greater than UNITSZ. If it is, VTAM uses more than one buffer to contain the message. The unused space in the last buffer is not allocated to the next message.

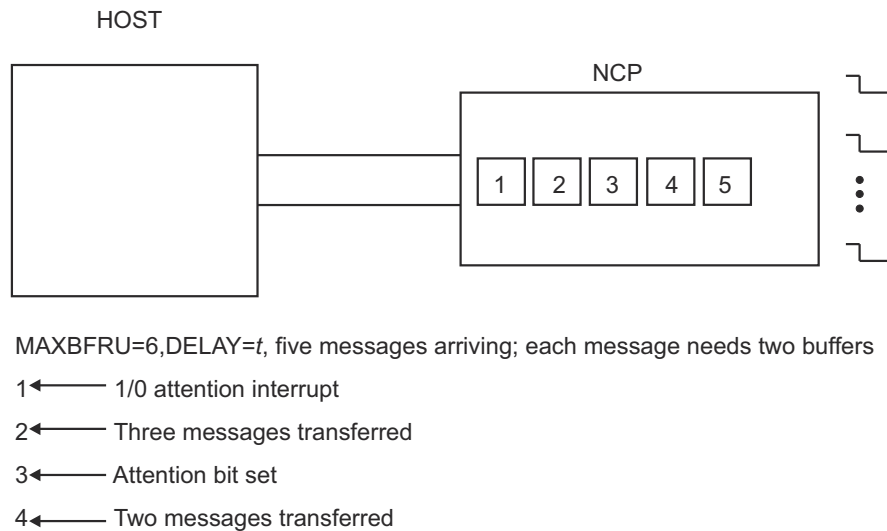


Figure 150. Multiple-buffer considerations

Figure 150 on page 553 assumes that MAXBFRU=6 and DELAY=t and that there are five messages arriving at the NCP before the time t expires. However, it also assumes that each message requires two I/O buffers. After three messages are received, the MAXBFRU limit is reached, and the NCP signals the host with an attention interrupt. The three messages are transferred, but two more remain. Therefore, the NCP signals VTAM that there is more data to read (RDATN tuning statistic). If the RDATN tuning statistic is high, NCP has more data to transfer to VTAM than VTAM has allocated buffers to process. Therefore, you can increase the MAXBFRU specification, thus allocating more VTAM buffers to reduce this condition.

However, coattailing has taken place; two messages have been transferred without having generated an attention interrupt. VTAM creates another channel program to read the remaining messages. Again, coattailing has occurred, because two messages have been transferred and only one attention interrupt has occurred. However, two I/O operations have been used where the data transfer might have been accomplished with only one. If UNITSZ were large enough to contain an entire message or if MAXBFRU were equal to 10, the extra I/O operation could have been avoided.

Guidelines for setting UNITSZ

To choose the best UNITSZ value, first determine the average size of inbound and outbound messages. Then, compare them to each other, and set the UNITSZ equal to the larger of the inbound or outbound message sizes. For channel-attached cluster controllers add 16.

Note: The message size is not determined by VTAM. You can determine the average message size by examining your application programs (for example, CICS and IMS). You can also use the NetView Performance Monitor to analyze the average message size for your network.

Suppose the inbound size is much smaller than the outbound size. Choose a UNITSZ at least large enough to contain one complete inbound message. However, the UNITSZ should be larger than the inbound message so that relatively few buffers are needed for the outbound message. Because the extra buffers increase path length, it is better not to specify too many buffers.

As another example, consider the case where the inbound message size is much greater than the outbound size. The process is similar to the previous one, except that the outbound message is the guide for choosing UNITSZ. Choose a UNITSZ that is greater than or equal to the outbound message size so that relatively few buffers are needed to contain one inbound message.

Where practical, choose UNITSZ so that no more than five to seven buffers are needed for transferring one message, whether it is inbound or outbound.

Coattailing for channel-to-channel operations

VTAM channel-to-channel I/O operations, which are similar to SNA controller operations, are impacted by factors such as VTAM I/O buffer size and the channel delay and maximum number of buffer pages on the CTC definitions. The following topics describe variables that can affect VTAM channel-to-channel performance.

Guidelines for setting DELAY

To select a DELAY operand value, use the following procedure:

1. During periods of the day when it is most likely that you are experiencing performance problems, turn on the tuning statistics. For channel-to-channel attachments, turn on tuning statistics at each processor using the following operator command:

```
F NET,TNSTAT,CNSL,TIME=1
```

2. For each recording of the tuning statistics, check the amount of coattailing that is occurring. For channel-to-channel connections, compare the TIMERS and CHNRM values of the channel-attachment major node associated with each host processor. The desirable TIMERS value is 0, but an occasional nonzero value is acceptable.
3. For channel-to-channel attachments, if the TIMERS value in any tuning statistics record is too large, deactivate the channel-attachment major node at each host processor, and activate a previously defined major node in which DELAY=0 has been specified on the LINE definition statement.

For other channel attachments, if the average number of PIUs seems too high (thereby indicating possible response time problems), you can activate alternate definitions with a lower DELAY value.

The following events cause PIUs on the queue to be sent before the DELAY time elapses:

- VTAM reaches the MAXBFRU limit.
- VTAM sends virtual route pacing responses.
- VTAM sends transmission priority two traffic.
- VTAM reaches the QDPTH value.

The default value for the DELAY operand is 0 (for example, no coattailing). Take the default under the following conditions:

- A low rate of transactions that use the channel-to-channel connection. A rate of under five transactions per second is considered low.
- Channel-to-channel traffic with frequent bursts of SENDs, such as from an application program that makes inquiries of the other host in rapid succession.

You can also analyze the average number of bytes transferred per I/O operation by dividing the total number of bytes transferred (RDBUF) by the number of READ channel programs (CHNRM).

Guidelines for setting I/O buffer size

Because data is transferred in the channel-to-channel buffers, increasing or decreasing the I/O buffer size has no effect on the channel program size.

Virtual route window sizes

The default minimum window size might be too small for a VTAM channel-to-channel virtual route. VTAM increases the window size only when the virtual route pacing response is not returned fast enough to prevent one window worth of PIUs from being queued up on the virtual route (a HELD condition). A channel-to-channel route is so fast that the virtual route pacing responses are turned around fast enough

to prevent HELD conditions from occurring very often. Because of this, the current virtual route window size tends not to increase, but to stay very close to the minimum VR window size.

The smaller the VR window size, the more virtual route pacing responses flow over the channel. This can cause higher CPU utilization than necessary.

Setting a higher minimum window size reduces the number of virtual route pacing responses and helps CPU utilization and throughput. For channel-to-channel one-hop routes, the default minimum VR window size is 1. The recommended value is 15. The recommended maximum virtual route pacing window size value is 50 for this environment.

If VTAM is expanding its IOBUF buffer pool to service a virtual route, the lack of VTAM buffers is considered major congestion for the route. Therefore, VTAM decreases the virtual route window size to the minimum for all VRs that it services. You should monitor this VTAM buffer pool to ensure that there is not excessive expansion and contraction if you are using dynamic buffering.

You can also use VTAM tuning statistics to analyze a virtual route impact on the channel-to-channel connection. A virtual route that uses transmission priority two and traverses the channel-to-channel connection causes VTAM to immediately schedule the data transfer operation. A virtual route pacing response is also high-priority traffic, which has the same effect. The PRI tuning statistic in VTAM indicates the number of times that a VTAM channel program is started to transfer this high priority data (TP2 or VR pacing response) to the other VTAM host. If this number is high and the channel-to-channel connection is not used extensively for TP2 traffic, the minimum virtual route pacing window size is probably too small. Also, the higher this number is in relation to the sum of TIMERS, QDPTH, and BUFCAP, the less outbound coattailing occurs.

For more information about pacing, see [“Session-level pacing”](#) on page 217.

Session-level pacing tuning considerations

Usage of the following storage areas affects the efficiency of pacing in the network:

- Input/Output buffer
- Application program data space
- Common storage manager

When storage usage hits a constraint level (70 percent usage) or critical level (90 percent usage), VTAM takes an appropriate action to reduce the storage usage. Constraint and critical levels are in part determined by user specifications and vary for each of the listed storage areas.

Input/output buffers

The IOBUF start option defines the amount of storage available for I/O buffers. The XPANLIM parameter is used to increase storage available for I/O buffers when reaching the constraint level.

For a detail description of the XPANLIM parameter, see [“Guidelines for dynamic expansion”](#) on page 541.

Application program data space

VTAM uses the value of the ASRCVLM operand coded on the APPL definition statement in calculating the constraint and critical levels for application program data space. Selecting a value for the ASRCVLM operand should reflect the high-usage periods of the data space storage.

To determine current usage value for the application program data space, issue a DISPLAY ID operator command, specifying the application program name. Message IST1633I contains the value specified for the ASRCVLM operand and message IST1634I contains the address space usage.

CSM storage

Threshold values for constraint and critical levels of CSM storage is defined by the IVTPRM00 parmlib member. See [z/OS Communications Server: New Function Summary](#) for additional information about specifying tuning values for CSM storage.

Appendix A. TSO/VTAM

TSO/VTAM provides the capability of using TSO through VTAM. TSO is a standard feature in MVS that provides conversational time sharing.

TSO/VTAM supports the following terminals:

- IBM 3270 Information Display System
- IBM 3290 Information Display System
- IBM 3767 Communication Terminal
- IBM 3770 Data Communication System
- IBM 5550 (as 3270) single-byte character set (SBCS) and double-byte character set (DBCS) modes
- IBM 8775 Display Terminal
- IBM Displaywriter

TSO/VTAM can also support the following non-SNA terminals in conjunction with the IBM Network Terminal Option licensed program:

- CPT-TWX Terminal
- IBM 2741 Communication Terminal
- World Trade Telegraph (WTTY)

The basic elements of TSO/VTAM are the terminal control address space (TCAS) and the VTAM terminal I/O coordinator (VTIOC). TCAS accepts logons from TSO/VTAM users and creates an address space for each user. VTIOC is the interface between TSO and VTAM; it coordinates data flow. TCAS and VTIOC, together with existing TSO components such as the LOGON scheduler, the application programs, the terminal monitor program, and the TSO service routines, make up a TSO/VTAM time-sharing system.

The user logon to the TCAS address space is passed to the TSO user address space without initially establishing a session between the TCAS address space and the TSO terminal. Instead, the session is established between the TSO terminal and the TSO user address space, avoiding session establishment and termination with the TCAS address space and avoiding your receiving message USSMSG7.

TSO/VTAM and TSO through VTAM can reside in the same host processor.

Note: z/OS limits the number of concurrent TSO/E sessions to eight.

Defining TSO to VTAM

To use TSO/VTAM with VTAM, define the following to VTAM:

- TCAS and each TSO user
- TSO/VTAM session parameters

For compatible logons, you also need to create and define to VTAM an interpret table.

Defining the TCAS application to VTAM

An APPL definition statement defines an application program to VTAM. Because TCAS and each TSO user are VTAM application programs, you must code APPL definition statements for them and put the definition statements in SYS1.VTAMLST.

You can reduce the number of operands you need to code on the APPL definition statements for each TSO user by using the GROUP definition statement under the APPL major node. All of the operands on the APPL definition statement except ACBNAME can be coded on the GROUP definition statement and allowed to sift down. You only need to code the unique name on the APPL definition statements for the TSO users if you use the GROUP definition statement to specify all the other operands.

Note: To avoid obtaining unnecessary CSA storage for TSO users, you should define TSO application programs to VTAM with EAS=1; do not use the default. TSO/VTAM obtains CSA storage independently of VTAM CSALIMIT.

Single-domain network

In a single-domain network in which there are no overriding network naming conventions, you can use the following technique to code the APPL definition statements for TCAS and for the TSO users:

- Code the following APPL definition statement for TCAS:

```
tsoa APPL ACBNAME=TSO,PRTCT=password,
      AUTH=(NOACQ,PASS,NVPACE,TSO,NOPO),EAS=1
```

The label *tsoa* is a unique network name. This name can be up to eight characters in length, but should be kept to less than eight so that it can be a prefix for the names on the subordinate TSO APPL definition statements.

- Code as many APPL definition statements, in the following format, as there will be users logged on to TSO/VTAM at one time:

```
TSOannnn APPL ACBNAME=TSOnnnn,PRTCT=password,
      AUTH=(NOACQ,PASS,NVPACE,TSO,NOPO),EAS=1
```

The label *tsoannnn* is a unique network name, which can be up to eight characters in length. It can be helpful to use the network name for TCAS as a prefix in the network name when coding APPL definition statements for each terminal logged on to the TSO/VTAM in this VTAM host concurrently. For the *nnnn* suffix, use sequential decimal integers starting with 0001. For the ACBNAME, the *nnnn* suffix is a decimal integer; the numbering must start with 0001 and be sequential.

Note: It is recommended that session pacing should not be used; NVPACE should be coded. However, in situations where pacing is necessary, code a high value, such as VPACING=20.

You need to use the same *password* for TCAS and each TSO/VTAM user. You must code a different application program name, in the form *TSOnnnn*, for each user. The *nnnn* suffix is a decimal integer; the numbering must start with 0001 and must be sequential. NOACQ and NOPO need not be coded; they are default values.

Multiple-domain network

For TSO sessions to be established in a multiple-domain network, you must code CDRSC definition statements for each TSO/VTAM application program unless you specify SSCPDYN=YES. Code the CDRSC definition statements in every VTAM host owning logical units that will log on to the application program.

In the VTAM host containing the TSO subsystem, it is not necessary to define the LUs in the other domains that issue logons to TSO/VTAM if dynamic CDRSC definition is authorized. Also, the manager in the domain of the SLU can be authorized to create CDRSCs for the TSO application program associated with the TSO user.

Code the following CDRM definition statements in each domain to avoid defining both the SLUs that log on to TSO/VTAM as CDRSCs in the TSO/VTAM domain and the TSO user application programs as CDRSCs in the SLU domain.

```
name1    CDRM   CDRDYN=YES,CDRSC=OPT,...
name2    CDRM   CDRDYN=YES,CDRSC=OPT,...
```

The CDRDYN operand determines whether a CDRM is authorized to dynamically create a CDRSC representing a cross-domain LU when a logon request is received from the LU (through an Initiate from the CDRM managing the LU).

The CDRSC operand determines whether the dynamic creation of CDRSC definitions is permissible when a logon request is received from the CDRSC manager identified by this CDRM definition statement.

In a multiple-domain network where TSO/VTAM is run in more than one domain, you must have special definition statements with unique names for TCAS and the terminals in your domain and interacting domains. Code the following APPL definition statement for TCAS in your domain:

```
tsoa      APPL  ACBNAME=TSO,PRCT=password,
              AUTH=(NOACQ,PASS,NVPACE,
              TSO,NOP0),EAS=1
```

The label *tsoa* is a unique network name. This name can be up to eight characters in length, but should be kept to less than eight so that it can be a prefix for the names on the subordinate TSO APPL definition statements.

Code as many APPL definition statements in the following format as the maximum number of sessions that will be established with TSO/VTAM in your domain at one time:

```
tsoannnn APPL  ACBNAME=TSOnnnn,PRCT=password,
              AUTH=(NOACQ,PASS,NVPACE,
              TSO,NOP0),EAS=1
```

The label *tsoannnn* is a unique network name, which can be up to eight characters in length. It can be helpful to use the network name for TCAS as a prefix in the network name when coding APPL definition statements for each terminal logged on to the TSO/VTAM in this VTAM host concurrently. For the *nnnn* suffix, use sequential decimal integers starting with 0001. For the ACBNAME, the *nnnn* suffix is a decimal integer; the numbering must start with 0001 and be sequential.

Code the following definition statement for each TCAS in another domain with which an LU in your domain communicates:

```
tsob      CDRSC  CDRM=name of VTAM manager for tsob
```

If a dynamic cross-domain resource definition is not authorized, code the following CDRSC definition statement in each domain that contains an SLU that can communicate with your domain, and code a CDRSC definition statement in the domain of TSO/VTAM for each cross-domain SLU that can log on to TSO/VTAM:

```
tsobnnnn CDRSC  CDRM=name of VTAM manager for tsobnnnn
```

tsob identifies the owning CDRM for each domain and must be included on each TCAS and CDRSC definition statement. The *nnnn* suffix is a decimal integer; the numbering must start with 0001 and be sequential.

Using model application program definitions for TSO/VTAM application programs

You can reduce the number of APPL definition statements for terminals that can be logged on to TSO/VTAM by using wildcard characters to define model application program definitions. See [“Model application program definitions” on page 278](#) for a full explanation of model application program definitions.

The following example shows how model application program definitions are used for TSO/VTAM application programs.

The following APPL definition statement defines TCAS.

```
tsoa      APPL  ACBNAME=TSO,PRCT=password,
              AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
```

The following APPL definition statement defines a model that is used to define the TSO address space as terminals log on to TSO/VTAM.

```
tsoa*     APPL  ACBNAME=TSO*,PRCT=password,MODSRCH=FIRST,
              AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1
```

Requirement: The label in these definitions must be 4 characters or less in length (without the *). In the previous examples, the labels *tsoa* and *tsoa** are 4 characters in length (without the *).

As each terminal logs on to TSO/VTAM, the * in TSO* on the ACBNAME operand and in the label *tsoa** will be replaced with the appropriate sequential 4-digit decimal integer; for example, 0001, 0002, and so forth. You do not need to code separate APPL definition statements for each terminal that can be logged on to TSO/VTAM concurrently. You need only the one model application program definition.

Notes:

1. If you plan to use model application program definitions for TSO/VTAM, then you must code MODSRCH=FIRST on the TSO application definition statement that represents the subordinate TSO APPLs.
2. With wildcard characters, you cannot limit the number of TSO address spaces created by limiting the number of application definitions you code. However, you can use the USERMAX operand in the TSOKEYxx parmlib member for this purpose.

Using BASENAME to define TSO/VTAM applications programs

The BASENAME operand, coded in the TSOKEYxx parmlib member, can be used to specify up to four characters for the "base" part of the TSO/VTAM application program name. The default value for BASENAME is TSO.

Using BASENAME, you are not limited to the names TSO0001, TSO0002, and so forth. For example, if you code

```
BASENAME=TTTT
```

the names will be TTTT0001, TTTT0002, and so forth.

The ACBNAME operand must be coded for TCAS and must specify TSO. The value you specify for BASENAME does not apply to the TCAS definition.

When using BASENAME, coding the ACBNAME operand on the APPL definition statements for TSO users is not necessary. The default value for the ACBNAME on an APPL definition statement is the name specified on the APPL definition statement. When BASENAME is used, TSO will use the same name for the APPL name and the ACBNAME. Therefore ACBNAME is not required on the TSO user APPL definition statement when BASENAME is used.

By specifying a unique BASENAME on every VTAM, you provide the name TCAS needs to start the TSO application address spaces, and provide the network-unique application name. Using the example above, you can code the following application program definitions:

```
tso1      APPL  ACBNAME=TSO,PRTCT=password,
               AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1

tttt0001  APPL  PRTCT=password,
               AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1

tttt0002  APPL  PRTCT=password,
               AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),EAS=1

tttt0003  APPL  PRTCT=password,
               AUTH=(NOACQ,PASS,NVPACE,TSO,NOP0),MODSRCH=FIRST,EAS=1

:
```

You can combine BASENAME with wildcard characters, eliminating the multiple application program definitions. The following example shows how to do this:

```
tso1      APPL  ACBNAME=TSO,PRTCT=password,
               AUTH=(NOACQ,PASS,NVPACE,
               TSO,NOP0),EAS=1

tttt*    APPL  PRTCT=password,
```

```
AUTH=(NOACQ,PASS,NVSPACE,TSO,NOP0),  
MODSRCH=FIRST, EAS=1
```

As each terminal logs on to TSO/VTAM, the * in the label tttt* will be replaced with the appropriate sequential 4-digit decimal integer; for example, 0001, 0002, and so forth.

Resource registration in an APPN network

To prevent incorrect session setups to TSO or TSO logon failures, the default resource registration is different for TSO applications than for other applications. Applications are registered to a directory server by default because they are likely to be the object of a search. Because multiple TSO applications can be defined using the same name, registering all TSO applications would cause misleading information to be stored in the directory server. Instead, TSO applications are registered according to the following conditions:

- If the APPL resource name is TSO and the ACBNAME is TSO, the default resource registration for this resource is NO registration.
- If the APPL resource name is tsoxxxxx and the ACBNAME is TSO, the default registration for this resource is CDSERVER (register to the central directory server).
- For the APPL resource names for the TSO address spaces (for example, the names ending in 0001, 0002, and so forth), the default is NO registration.

It is assumed that AUTH=TSO is coded on each APPL definition statement.

Note: It is recommended that you not modify the registration of TSO applications from the default to avoid session setup or logon problems.

Defining TSO/VTAM session parameters

VTAM needs to know the session protocols required by the terminals from which individual TSO users log on. This information is provided through logon mode tables and through VTAM and NCP definition statements.

Logon mode tables associate session protocols with device types. IBM provides a default logon mode table that describes session protocols for common devices and situations. However, you might need to replace or supplement the default table.

In addition, a specific Class of Service name can be chosen. The Class of Service name used for the session is selected from the logon mode entry. If you omit the COS operand on the MODEENT macroinstruction, the default COS is chosen. To choose a specific COS, you need to include the COS name on the MODEENT macroinstruction as follows:

```
MODEENT COS=cos name,logon mode parameters
```

The connection between a logon mode entry and a device is made using VTAM operands on the LU, LOCAL, or TERMINAL definition statements during VTAM and NCP definition. Code the DLOGMOD operand to identify the name of a logon mode table entry containing the session parameters used for the 3270 device being defined. Code the optional MODETAB operand to identify the logon mode table containing the entry. If you do not code MODETAB, the IBM-supplied logon mode table (ISTINCLM) is used.

You can also code the MODETAB and DLOGMOD operands on the GROUP, LINE, PU, LU, or CLUSTER definition statements to take advantage of the sift-down effect.

TSO/VTAM support of 3270 devices

TSO/VTAM supports the following two types of 3270 devices:

- Non-SNA devices attached locally or by a bisynchronous line protocol
- SNA devices attached by SDLC links

If no logon mode table entry exists for a 3270 device, TSO/VTAM assumes that the device is non-SNA (LU0) and uses the buffer size in the SCRSIZE parameter of the TSOKEY00 parmlib member. The following are valid values for the SCRSIZE parameter:

- 480 [12x40]
- 1920 [24x80]
- 2560 [32x80]
- 3440 [43x80]
- 3564 [27x132]
- 9920 [62x160]

For non-SNA devices with different screen sizes or special features and for SNA devices (LU2), you need to provide a logon mode table entry.

Use the MODETAB and MODEENT macroinstructions to define 3270 characteristics in logon mode tables and their entries. The PSERVIC operand of the MODEENT macroinstruction defines device LU type, buffer sizes, single-byte language, and QUERY capability (that is, programmed symbol sets, extended color, extended data stream, or extended highlight support). The LANG operand of the MODEENT macroinstruction defines the single-byte language of the device and whether it should be queried for double-byte capability.

Following are examples of definitions of MODEENT table entries.

MODEENT macroinstruction for non-SNA 3270 devices

The FMPROF, TSPROF, PRIPROT, SECPROT, and COMPROT values shown in the following example are the same as those used in the IBM-supplied logon mode table, ISTINCLM.

```
name      MODEENT FMPROF=X'02',
          TSPROF=X'02',
          PRIPROT=X'71',
          SECPROT=X'40',
          COMPROT=X'2000',
          PSERVIC=X'...1'
```

Note:¹ — See [Table 63](#) on page 563.

MODEENT macroinstruction for SNA 3270 devices

The RUSIZES operand in the following example defines a 256-byte maximum secondary logical unit RU send size and a 1024-byte maximum primary logical unit RU send size. If the maximum primary logical unit RU size is 0, TSO/VTAM defaults to 1024 bytes.

```
name      MODEENT FMPROF=X'03',
          TSPROF=X'03',
          PRIPROT=X'B1',
          SECPROT=X'90',
          COMPROT=X'3080',
          RUSIZES=X'8587',
          PSERVIC=X'...1'
```

PSERVIC operand of the MODEENT macroinstruction

Associated with each terminal is a default logon mode table entry. Within this entry, the primary and alternate screen sizes are specified, and a code that describes how these screen sizes are used.

In the MODEENT macroinstruction, the 11th byte of the PSERVIC can have the value of X'03' (PSERVIC=X'.....03..'). When this field is specified as X'03', the primary screen size is 24x80, and the alternate screen size is determined by the device. TSO/VTAM queries the device to determine the alternate screen size.

In the MODEENT macroinstruction, the 2nd and 11th bytes of the PSERVIC can have the value of (PSERVIC=X'..80.....00..') for a device with extended data stream capability. When these fields are

specified, the primary screen size and the alternate screen size are determined by the device. TSO/VTAM queries the device to determine the primary and the alternate screen sizes.

You do not have to code different logon mode table entries for TSO users when the only difference is the screen size. Users logging on to TSO with a screen size different from the default do not have to specify a logon mode table entry. The screen size is determined dynamically.

Note: The 3274 Terminal Controller must be Release 65, or higher, or the logon fails.

Code the 12 bytes of device-specific hexadecimal data of the PSERVIC operand as described in [Table 63](#) on page 563.

<i>Table 63. Coding the device-specific hexadecimal data of PSERVIC</i>	
Hexadecimal code	Related device
X'00..00000000.....00'	For non-SNA (LU0)
X'02..00000000.....00'	For SNA (LU2)
X'..00.....'	Device without extended data stream capability
X'..80.....'	Device with extended data stream capability
X'.....0000000001..'	Buffer size 480 only (12x40)
X'.....0000000002..'	Buffer size 1920 only (24x80)
X'.....0C280C507F..'	Buffer sizes 480 or 960 (12x40 or 12x80)
X'.....185020507F..'	Buffer sizes 1920 or 2560 (24x80 or 32x80)
X'.....18502B507F..'	For buffer sizes 1920 or 3440 (24x80 or 43x80)
X'.....18501B847F..'	For buffer sizes 1920 or 3564 (24x80 or 27x132)
X'.....18503EA07F..'	Buffer size 9920 only (62x160)
X'.....0000000003..'	Primary buffer size 1920 only (24x80). If a device is extended data stream capable (see above), TSO/VTAM queries the device to determine the alternate size.

To prevent switching of screen sizes on a device that has more than one size possible, code the screen size you want in the primary area of PSERVIC and a X'7E', which means no switching. Following is an example:

PSERVIC=X'.....1B8400007E.. ' For 3564 buffer only.

For details on coding the PSERVIC operand for your particular device, see your component description.

For details of the bit settings in the PSERVIC operand that represent bytes 13 - 24 of the session parameters, see [z/OS Communications Server: SNA Programming](#).

LANG operand of the MODEENT macroinstruction

For TSO/VTAM, this operand affects (but does not necessarily decide) the following functions:

- The language of certain TSO/VTAM terminal user messages
- The language returned on the GTTERM macro
- The single-byte character set (SBCS) filter
- The device double-byte character set (DBCS) capability

The high-order bit of LANG is used, in conjunction with other LOGMODE information, to determine if a device is queried. A device is queried if it is extended data stream capable and either the high-order bit of LANG (from the MODEENT macroinstruction) is on, or the LOGMODE indicates to query for alternate screen size. The query of the device can affect each of the functions listed above.

Note: The high-order bit of the LANG operand is always obtained from the logon mode table entry (specified by the LOGMODE operand). If the terminal user enters a LOGON command, the LANG or LANGTAB operands of the LOGON command will override the low-order 7 bits specified by the LANG operand of the MODEENT macroinstruction.

The language of TSO/VTAM terminal operator messages IKT00201I through IKT00204I (logon failure messages)

TSO/VTAM attempts to retrieve these messages from the MVS message service (MMS), so that they can be issued in the language required by the terminal user. The language for all other TSO/VTAM terminal user messages is set (determined) by the PLANG operand of the PROFILE command.

If the device is queried and the high-order bit of LANG is on, the language obtained from the query (if known) is used for the language of the messages.

The language returned on the GTTERM macro

The TSO terminal monitor program (TMP) or command processor may issue the GTTERM macro to determine the language of a TSO/VTAM user. The PROFILE command does not affect the results of the GTTERM macro.

If the device is queried and the high-order bit of LANG is on, the language obtained from the query (if known) is used for the language returned by the GTTERM macro.

The single-byte character set (SBCS) filter

If the device is queried, the appropriate SBCS filter is used (U.S. English filter, Katakana filter, or no filter). If the device is not queried, the SBCS filter is determined from the language value.

TSO/VTAM uses this filter in TPUT data validation and translation. TSO/VTAM performs character filtering for the FULLSCR, EDIT, ASIS, and NOEDIT options of the TPUT and TGET macros.

The DBCS capability of a device

If the query reply indicates that a device is DBCS capable, TSO/VTAM enables mixed DBCS strings to be sent and received from the device. If not queried, the device is assumed incapable of DBCS processing.

If you want to disable the character filtering of TSO/VTAM, code LANG=X'7F' for a device that does not support double-byte capability, and code LANG=X'FF' for a potential double-byte character set (DBCS) device.

If an incorrect LANG operand value is specified or if LANG is not specified at all, the single-byte language defaults to U.S. English.

The low-order 7 bits indicate the single-byte language that is to be used. If a query is issued to determine double-byte capability and the high-order bit of LANG is on, any supported single-byte language value returned in the query reply overrides the value specified on the LANG operand.

The following indicates the bit settings for single-byte languages:

Bit

Description

0

Indicates that the device should not be queried for language information.

1

Indicates that the device should be queried for language information. This value should be used for devices that will potentially process DBCS data.

. nnn nnnn

Indicates the language. For a list of valid values, see [z/OS Communications Server: SNA Programming](#).

. 111 1111

Indicates that SBCS filtering is disabled (no SBCS filter). Although not a valid language, this value can be used to disable character filtering. Note that any language other than U.S. English (X'00' or X'01') or Katakana (X'11') will also disable SBCS filtering.

Defining the 3790/3270 configuration to TSO/VTAM

The 3270 attached to a 3790 (Version 7) uses the LU type 2 protocol. You need to give session parameters to TSO/VTAM through a logon mode table entry to properly identify the model number or screen sizes of the device being used. The 3270 operator selects FPID 932 when logging on to the 3790 to activate the 3270 data stream compatibility function. During the logon procedure, the operator is prompted for the application program ID and the logon mode table entry containing the session parameters to be used for that session. The information is transmitted to VTAM by the Initiate Self request generated by the 3790. If the operator does not supply a logon mode table entry, the 3790 uses the standard logon mode entry name EMU3790.

So that the terminal operator does not have to know what logon mode entry name to supply, create a separate logon mode table for each 3270 model attached to the 3790, and include the default entry name EMU3790 in each table with the correct session parameters for that device.

Use the MODETAB and MODEENT macroinstructions to define the logon mode table and its entry. The PSERVIC operand on the MODEENT macroinstruction carries the model number or screen sizes in row and column form. Following is an example of defining a logon mode entry for a 3270 attached to a 3790 with a screen size of 1920 bytes (Model 2).

MODETAB macroinstruction

```
LU2TABLE MODETAB
```

MODEENT macroinstruction

```
EMU3790  MODEENT  LOGMODE=EMU3790,
                  FMPROF=X'03',
                  TSPROF=X'03',
                  PRIPROT=X'B1',
                  SECPROT=X'90',
                  COMPROT=X'3080',
                  PSERVIC=X'020000000000000000000200'
```

For a description of the PSERVIC values, see [“PSERVIC operand of the MODEENT macroinstruction” on page 562](#).

The association between the logon mode entry and the device is made on the LU definition statement during network definition to VTAM. Use the optional MODETAB operand to identify the logon mode table containing the entry. If MODETAB is not coded, the IBM-supplied logon mode table (ISTINCLM) is used.

Defining 2741, TWX, or WTTY terminals to TSO/VTAM

A 2741, TWX, or WTTY device attached to a communication controller can be used with TSO/VTAM through its LU type 1 protocol managers. These devices are identified to VTAM during the NCP generation process. (For details, see *Network Terminal Option Migration and Resource Definition* and the *NCP, SSP, and EP Resource Definition Guide*.) Because the NCP translates ASCII line code to EBCDIC for these devices, make sure that any logon mode entry named by the DLOGMOD and MODETAB operands on the PU or LU definition statements identifies the device to VTAM as EBCDIC. Following is an example of an ASCII TWX logon mode table entry definition that has the alternate code indicator in the COMPROT field set off (to indicate EBCDIC to TSO/VTAM).

MODETAB macroinstruction

```
TWXTABLE MODETAB
```

MODEENT macroinstruction

```
TWXDEVIC MODEENT LOGMODE=TWXDEVIC,
                  FMPROF=X'03',
                  TSPROF=X'03',
                  PRIPROT=X'B1',
                  SECPROT=X'90',
```

```
COMPROT=X'3040'  
DCODE=X'00'
```

The DCODE operand on the MODEENT macroinstruction indicates to TSO/VTAM whether a TWX device is a keyboard display or a keyboard printer. If you specify DCODE to indicate that the device is a keyboard display, TSO/VTAM inhibits the display of the password when the password is entered at the device.

Defining an interpret table for compatible logons

Because TSO/VTAM uses VTAM logon facilities, the TSO LOGON command is not supported. You can make this not apparent to terminal users by defining an interpret table to allow logon requests to have the same format as that used by the TSO LOGON command. The following interpret table definition can be used:

```
INTBL    INTAB  
         LOGCHAR APPLID=(APPLICID,TSO),SEQNCE='LOGON'  
         LOGCHAR APPLID=(APPLICID,TSO),SEQNCE='logon'  
         ENDINTAB  
         END
```

This interpret table allows a user to enter a logon command in uppercase or lowercase letters.

Defining TSO to MVS

This section describes how to define TSO to MVS.

Writing a procedure to start TSO/VTAM time sharing

You need to write a cataloged procedure to start TSO/VTAM time sharing and include it in either SYS1.PROCLIB or your own procedure library. For details, see [z/OS TSO/E Customization](#).

Creating a TSOKEY00 PARMLIB member

TSO/VTAM time sharing parameters determine how time sharing buffers should be controlled, the maximum number of users, and other options. Unless you want the IBM default parameters to apply, you must provide a parmlib member containing the parameters that you want. Information about initialization and tuning for your operating system describes the parameters you can code, how to code them, and the IBM defaults.

You can use MVS system symbols in the TSOKEY00 parmlib member. For information about using MVS system symbols, see [“Using MVS system symbols” on page 30](#).

Defining TCAS program properties

TCAS must have an entry in the MVS program properties table (PPT). This entry designates the execution attributes of TCAS, which should be as follows:

- The program cannot be canceled.
- Unique protection key 6 must be assigned to the program.
- The program is privileged and is not swapped unless the address space is in a long wait.
- The program is a system task and is not timed.
- There is no host affinity.

For coding conventions, see initialization and tuning information for your operating system.

Implementing TSO/VTAM

This section has the following subsections:

- Translation tables
- Coding TSO/VTAM exit routines

- Security
- Performance
- 3270 large screen considerations
- TSO considerations
- Multicultural support for TSO user messages
- Operating VTAM under TSO

Translation tables

Translation tables allow TSO/VTAM users to replace internally characters that are not available on a keyboard. If you call for character translation, translation tables (either your own or those supplied by IBM) are used. For more information, see [z/OS TSO/E Customization](#).

Coding TSO/VTAM exit routines

Exit routines provide the following functions needed in TSO/VTAM:

- Input and output editing that replaces or supplements IBM-supplied editing
- Attention handling that replaces IBM-supplied attention handling
- Support for terminals not supported by TSO/VTAM
- Verifying whether a terminal is capable of receiving a TSO/VTAM user message in the language selected by the TSO user

For information about coding TSO/VTAM exit routines, see [z/OS Communications Server: SNA Customization](#).

Security

The TSO subsystem is considered a secure application program. That is, confidential data is handled on behalf of the user in ways that prevent unauthorized disclosure of the data. The CONFTXT parameter in the TSOKEY00 parmlib member determines whether output data is considered confidential text. The default is that the data is considered confidential. TSO/VTAM protects tracing of user data by setting the CONFTXT indicator in the NIB at the time the user logs on.

If CONFTXT=NO, VTAM can perform buffer or I/O traces on the data. If CONFTXT=YES (the default), the data is considered confidential and no data is recorded. The CONFTXT parameter, however, does not apply to the TSO type VTAM trace for TPUT/TPG/TGET buffers; these are always traceable. For details, see information about initialization and tuning for your operating system.

VTAM supports TSO message security by invoking RACF services to provide resource access control for:

- Cross-address space TPUTs (such as the TSO SEND command), which:
 - Control who can send messages to whom
 - Ensure that a message will be received by the intended user
 - Ensure that a cross-address-space message can be received only by a user with a security classification that is equal to or greater than the sender
- Requests to open an ACB from a non-APF authorized application program or processor.

Note: The installation must ensure that both the sender and receiver of TSO/VTAM messages are authorized with the proper security level in the security management product. The TSO/VTAM user IDs should be registered with a class of SMESAGE in the security management product.

Performance

Following are suggestions for improving TSO/VTAM performance:

- Terminal users should stack input data whenever possible. This means using the multiple-line input technique on 3270 terminals and buffered SDLC transmission on 3767 and 3770 terminals. Stacking input results in a decreased number of data transmissions to TSO. Users of 3270 terminals can stack up to a full screen of data; 3767 and 3770 terminal users can stack from 256 bytes to 2000 bytes of data, depending on the terminal buffer size.
- Users of 3767 and 3770 terminals can reduce idle time by "typing ahead." This means that whenever the terminal is not receiving data or is transmitting data, the user can enter input.
- Users of 3276 and 3278 terminals can expedite unlocking the keyboard if you use BREAK mode. BREAK mode reduces idle time by allowing typing ahead.
- Users of 3270 terminals can shorten the TSO/VTAM processing path when handling input if you use automatic line numbering. This reduces system overhead and expedites unlocking the keyboard.

3270 large screen considerations

The following information describes screen management techniques for IBM 3270 terminals that have default and alternate screen sizes determined by the logon mode table entries used for them at logon. The default size is accessed by sending an Erase/Write (EW) command to the terminal, and the alternate size is accessed by sending an Erase/Write Alternate (EWA) command.

TSO/VTAM screen management

TSO/VTAM can manage the screen for TSO application programs that use line-oriented I/O. In this case, TSO/VTAM uses the maximum screen size for the device, whether it is the default or the alternate in the logon mode table entry. When screen erasure is necessary, TSO/VTAM uses either the EW or the EWA command as required to access the larger screen size.

Full-screen application program screen management

An IBM-supplied or user-supplied application program that uses full-screen I/O in TSO full-screen mode can be used to manage the screen. TSO/VTAM uses the screen size used by the application program. If the CLEAR key is pressed, TSO/VTAM erases the screen with either an Erase/Write (EW) or Erase/Write Alternate (EWA) command depending on which was last used by the application program.

TSO considerations

Information about the TSO-related steps that must be performed before a TSO/VTAM terminal user logs on is in [z/OS TSO/E Customization](#) and in information about initialization and tuning for your operating system. (Make sure you have updated these books with TSO/VTAM supplements before referring to them.) The following list outlines the steps and refers you to the appropriate publication:

- Write LOGON cataloged procedures, and include them in SYS1.PROCLIB. See [z/OS TSO/E Customization](#).
- Construct the TSOKEY00 member or an alternate parmlib member (or an alternate data set) to set VTIOC parameters. See information about initialization and tuning for your operating system.
- Include SYS1.CMDLIB in a LNKSTxx parmlib member or in a LOGON cataloged procedure. See information about initialization and tuning for your operating system.
- Create or convert the user attribute data set (UADS) and the broadcast data set. See [z/OS TSO/E Customization](#).
- Build translation tables if you want character translation. See [z/OS TSO/E Customization](#).
- Write the procedure that starts TSO/VTAM time sharing. See [z/OS TSO/E Customization](#).
- Write any command exit routines you plan to use. See [z/OS TSO/E Customization](#).

Note: You are strongly encouraged not to VARY LOGON to TSO or code TSO as a LOGAPPL because unpredictable results can occur. For instance, if the original BIND fails, TSO messages and VTAM

messages may be repeatedly issued to the console indicating the session setup failure. In some cases this may not stop even if the VARY NOLOGON command is used to break the automatic logon relationship.

Multicultural support for TSO/VTAM user messages

You can tailor TSO/VTAM user messages using the MVS Message Service (MMS). You define TSO/VTAM user messages to MMS the same way you define USS user messages. For details, see [“Defining USS messages to the MVS message service” on page 200](#).

The TSO user selects the language used for TSO/VTAM messages using the PROFILE command with the PLANG operand. TSO/VTAM then retrieves its messages from MMS using the language specified by the user. This language continues to be used until the TSO/user changes the language (by reissuing the PROFILE command).

Before the TSO user can issue the PROFILE command, TCAS may issue messages IKT00201I through IKT00204I. The language used for these messages is specified on either the LANG operand of the MODEENT macroinstruction or on the USS logon command using the LANG or LANGTAB operands. The language is passed from these sources to the TSO/VTAM logon exit.

For information about how MMS translates messages, see [Appendix G, “Message translation using the MVS Message Service,” on page 609](#).

Operating VTAM under TSO

A TSO operator is defined as any TSO user with the authority to use the CONSOLE command. MVS or subsystem operators running under TSO are treated as extended multiple consoles. Authorized users can establish a console session with MVS console services. After the session is active, users can enter MVS and subsystem commands and obtain solicited and unsolicited messages.

By representing all consoles and operators as a 4-byte console ID, multiple console support (MCS) can extend its current command and message services to TSO operators. This allows TSO operators to send VTAM commands and receive VTAM messages.

To enable this function, you need to grant individual users CONSOLE command authority, so that they can use the MVS CONSOLE command. There are three ways that you can give CONSOLE authority to users:

- Using RACF facilities (RACF RDEFINE and RACF PERMIT commands)
- Using the TSO logon preprompt exit routine
- Using the TSO CONSOLE exit routine

For information about giving CONSOLE authority, the levels of authority, and the CONSOLE command, see [z/OS TSO/E Customization](#).

A TSO operator can use the MVS Message Service (MMS) to retrieve translated messages, including VTAM, logon manager, and TSO/VTAM messages for network operators. The TSO operator can select a language using the PROFILE command with the PLANG operand. For information about how MMS translates messages, see [Appendix G, “Message translation using the MVS Message Service,” on page 609](#). For details on defining translated network operator messages, see the [z/OS Communications Server: SNA Resource Definition Reference](#).

Appendix B. Storage estimate worksheets

This appendix describes how to estimate the virtual storage required to run z/OS Communications Server on the z/OS operating system. From the following list, select and review the applicable worksheets:

- APPN
 - Interchange node (ICN) or network node (NN)
 - Migration data host (MDH) and end node (EN)
- Subarea
 - Communication management configuration (CMC)
 - Data host (DH)

This appendix also contains an APPL EAS storage estimate table; see [“APPL EAS storage estimates” on page 579](#).

Users of APPC will notice an increase in storage utilization because VTAM will now allocate an additional 160 bytes (for a mini-VIT trace) per control block representing a single APPC conversation. See [z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT](#) for information about VTAM Internal Trace (VIT) tracing.

Users of APPN will notice an increase in storage utilization because VTAM will now allocate additional storage for TRS (Topology and Routing Services) topology traces. To calculate the increase, add the following:

- For the TRS topology trace where topology record deletions are recorded, one to ten 4K pages of storage will be allocated. One 4K page will be allocated at VTAM initialization. When that page is full of trace entries, another page will be allocated, up to a maximum of ten 4K pages.
- For the NDREC (node record) traces, 110 bytes of storage will be allocated for each node in the network or sub-network.
 - In a network node, this would include:
 - All network nodes
 - All served end nodes (the network node has acted as the NNS or DLUS for the end node)
 - All virtual nodes
 - In an end node, this would include:
 - Network nodes with which the end node has established connections
 - End nodes with which the end node has established connections
 - Virtual nodes through which the end node has established connections
- For the TGREC (TG record) traces, 180 bytes of storage will be allocated for each TG in the network or subnetwork.
 - In a network node, this would include:
 - Two TGs for every link between network nodes (one TGREC represents the connection in one direction and one TGREC represents the connection in the reverse direction)
 - One TG for every connection that a VTAM end node served by this network node server has with any other network node, end node, or virtual node
 - One TG for every connection that a DLUR end node served by this DLUS has with any other network node, end node, or virtual node
 - Two TGs for every connection between network nodes and virtual nodes (one TGREC represents the connection in one direction and one TGREC represents the connection in the reverse direction)
 - One TG for every connection to an adjacent end node or an adjacent DLUR end node

- In an end node, this would include:
 - One TG for every link from this end node to any other network node, end node, or virtual node.

General information

These worksheets address z/OS Communications Server storage above 16MB; storage below 16MB is allocated as 150KB common storage, and 64KB private storage.

The heading "DISPLAY STATS ID" refers to the particular statistic identifier issued by the D NET, STATS command; this statistic identifier is used in the corresponding step of the worksheet. There are some statistics that D NET, STATS does not capture; they are indicated by N/A.

"Dynamic storage" refers to storage created in response to a need, and required only so long as the process using it stays active. Dynamic storage can be used to establish normal sessions, and for error recovery. Dynamic storage usage varies by configuration; it is related to the number of sessions being established. The worksheet yields an approximation of dynamic storage needed for a given configuration, assuming worst case system recovery.

Estimation of z/OS Communications Server storage is based on the following assumptions. (These assumptions reflect no implied or expressed recommendation.)

- All PUs and LUs defined in the configuration are active.
- Tables and user exit routines are not used.
- Default buffer pool values, dynamic buffer expansion, and default start options are used.
- Dynamic storage requirements are based on full system recovery.

In order to establish storage estimates, try the following steps:

1. Set REGION=0 for Private Storage, and CSALIMIT=0 for CSA/ECSA.
2. Monitor storage for approximately 2 or 3 weeks to determine what the peak values are.
3. When peak value is identified, then establish values for REGION and CSALIMIT. Use 1.5 times the peak value noted during the 2 or 3 week observation period.

APPN interchange node or network node

For APPN interchange node (ICN) or network node (NN) configurations, use the following worksheet:

Table 64. Worksheet for APPN interchange node or network node storage					
Step	Description	Formulas (decimal)	Common	Private	DISPLAY STATS ID
1	Number of 4KB pages allocated for the SNA internal trace table	(trace tab + 4) * 4096	= _____		2
2	Determine the size of the IOBUF parameter. This value is defined in the SNA start list. These buffers hold data transmitted to and from SNA.	Used in questions 4 & 12			3
3	Number of channel-attached communication controllers (NCPs) activated and owned by this SNA.	COMMON STORAGE: (NCPs * 1200)	= _____		5

Table 64. Worksheet for APPN interchange node or network node storage (continued)

Step	Description	Formulas (decimal)	Common	Private	DISPLAY STATS ID
4	Sum value of MAXBFRU parameters for all channel-attached communications controllers activated by this SNA. MAXBFRU is defined in the HOST definition statement of the NCP channel-attached major node definition.	COMMON STORAGE: ((IOBUF size + 98) * 2 * maxbfu NCPs)	= _____		3 & 6
5	Number of PUs defined in this SNA. Include all PUs defined to SNA in PU definition statements, and controllers defined in CLUSTER definition statements. Include locally attached, remotely attached, dynamically added, switched, ICA and NCP (including NTRI) PUs.	PRIVATE STORAGE: (defined PUs * 1000)		= _____	48 & 67
6	Number of device type LUs defined in this SNA. Include the locally attached LUs, and LUs attached through an NCP. In addition, include LUs that are defined to SNA in LU definition statements, and those devices defined in TERMINAL statements. Do not include applications.	PRIVATE STORAGE: (defined LUs * 820)		= _____	50
7	Number of independent LUs defined locally, remotely or by way of CDRSC. Include all independent LUs for which SNA provides boundary function services, and all NTRI independent LUs.	COMMON STORAGE: (indep LUs * 270) PRIVATE STORAGE: (indep LUs * 400)	= _____	= _____	80
8	Number of LU 6.2 sessions with application LUs owned by this SNA. LU 6.2 sessions are valid only for applications where APPC=YES is specified in the APPL major node definition. Include all same domain, cross domain, and cross network LU 6.2 sessions.	COMMON STORAGE: (LU6.2 ses * 840)	= _____		58

Table 64. Worksheet for APPN interchange node or network node storage (continued)					
Step	Description	Formulas (decimal)	Common	Private	DISPLAY STATS ID
9	Number of device type LUs owned by this SNA, and in session with an application program owned by this SNA.	COMMON STORAGE: (LUs w/appls * 500) PRIVATE STORAGE: (LUs w/appls * 420)	= _____	= _____	71
10	Number of device type LUs owned by this SNA, and in session with an application program owned by another SNA.	PRIVATE STORAGE: (cross node LU ses * 420)		= _____	73
11	Number of LU 6.2 sessions with both LUs owned by this SNA.	PRIVATE STORAGE: (LU6.2 ses * 450)		= _____	77
12	Number of device type (nonapplication) LUs in session with a TSO application.	COMMON STORAGE: (TSO LUs) * (2300 + IOBUF size)	= _____		55 & 3
13	Number of ENs that establish CP-CP sessions with this SNA.	PRIVATE STORAGE: (adj end node * 3170) COMMON STORAGE: (adj end node * 920)	= _____	= _____	104
14	Number of transmission groups used between this node and attached, or served, end nodes.	PRIVATE STORAGE: (end node TGs * 690)		= _____	142
15	Number of transmission groups used between this node and other network nodes.	PRIVATE STORAGE: (network node TGs * 690)		= _____	143
		Total Common = _____ 1024	= _____	KB (totcom)	
		Total Private = _____ 1024		= _____	KB (totpri)

Table 65. Summary of worksheet, APPN interchange node or network node storage		
ICN/NN configuration description	Common	Private
Calculated COMMON storage for ICN/NN configuration above	= _____KB (totcom)	
DYNAMIC COMMON STORAGE (4 x totcom)	= _____KB	
SNA SYSPLEX STORAGE (if used)	+100KB	
SNA BASE STORAGE (COMMON)	+ 3002KB	

Table 65. Summary of worksheet, APPN interchange node or network node storage (continued)		
ICN/NN configuration description	Common	Private
TOTAL SNA COMMON STORAGE	= _____KB	
Calculated PRIVATE storage for ICN/NN configuration above		= _____KB (totpri)
DYNAMIC PRIVATE STORAGE (2 x totpri)		= _____KB
SNA SYSPLEX STORAGE (if used)	+750KB	
SNA BASE STORAGE (PRIVATE)		+ 7057KB
TOTAL SNA PRIVATE STORAGE		= _____KB

APPN migration data host and end node

For an APPN migration data host (MDH) or end node (EN) configuration, use the following worksheet:

Table 66. Worksheet for APPN migration data host and end node					
Step	Description	Formulas (decimal)	Common	Private	DISPLAY STATS ID
1	Number of 4KB pages allocated for the SNA internal trace table.	(trace tab + 4) * 4096	= _____		2
2	Number of device-type LUs owned by this SNA in session with an application program owned by this SNA.	PRIVATE: (LUs w/appls * 420) COMMON STORAGE: (LUs w/appls * 500)	= _____	= _____	71
3	Number of cross node sessions between an application program in this SNA and a device type LU owned by another node or SNA.	PRIVATE: (cross node appl * 980) COMMON STORAGE: (cross node appl * 540)	= _____	= _____	112
Total Common = _____ 1024			= _____KB (totcom)		
Total Private = _____ 1024				= _____KB (totpri)	

Table 67. Summary of APPN migration data host and end node		
MDH/EN configuration description	Common	Private
Calculated COMMON storage for MDH/EN configuration from above	= _____KB (totcom)	
DYNAMIC COMMON STORAGE (2 x totcom)	= _____KB	
SNA SYSPLEX STORAGE (if used)	+100KB	

Table 67. Summary of APPN migration data host and end node (continued)		
MDH/EN configuration description	Common	Private
SNA BASE STORAGE (COMMON)	+ 3002KB	
TOTAL SNA COMMON STORAGE	= _____KB	
Calculated PRIVATE storage for MDH/EN configuration from above		= _____KB (<i>totpri</i>)
DYNAMIC PRIVATE STORAGE (2 x <i>totpri</i>)		= _____KB
SNA SYSPLEX STORAGE (if used)	+750KB	
SNA BASE STORAGE (PRIVATE)		+ 7057KB
TOTAL SNA PRIVATE STORAGE		= _____KB

Subarea data host

For a subarea data host (DH) configuration, use the following worksheet:

Table 68. Worksheet for subarea data host					
Step	Description	Formulas (decimal)	Common	Private	DISPLAY STATS ID
1	Number of 4KB pages allocated for the SNA internal trace table.	(trace tab + 4) * 4096	= _____		2
2	Number of device type LUs owned by this SNA in session with an application program owned by this SNA.	PRIVATE: (LUs w/appls * 420) COMMON STORAGE: (LUs w/appls * 500)	= _____	= _____	71
3	Number of cross domain sessions between an application program in this SNA and a device type LU owned by another SNA.	PRIVATE: (cross node appl * 910) COMMON STORAGE: (cross node appl * 540)	= _____	= _____	112
		Total Common = _____ 1024	= _____KB (<i>totcom</i>)		
		Total Private = _____ 1024		= _____KB (<i>totpri</i>)	

Table 69. Summary of subarea data host		
DH configuration description	Common	Private
Calculated COMMON storage for DH configuration from above	= _____KB (<i>totcom</i>)	
DYNAMIC COMMON STORAGE (2 x <i>totcom</i>)	= _____KB	
SNA BASE STORAGE (COMMON)	+2953KB	
TOTAL SNA COMMON STORAGE	= _____KB	

Table 69. Summary of subarea data host (continued)		
DH configuration description	Common	Private
Calculated PRIVATE storage for DH configuration from above		= _____KB (totpri)
DYNAMIC PRIVATE STORAGE (2 x totpri)		= _____KB
SNA SYSPLEX STORAGE (if used)	+150KB	
SNA BASE STORAGE (PRIVATE)		+ 5507KB
TOTAL SNA PRIVATE STORAGE		= _____KB

Subarea communication management configuration

For a subarea communication management configuration (CMC), use the following worksheet:

Table 70. Worksheet for subarea communication management configuration					
Step	Description	Formulas (decimal)	Common	Private	DISPLAY STATS ID
1	Number of 4KB pages allocated for the SNA internal trace table.	(trace tab + 4) * 4096	= _____		2
2	Determine the size of the IOBUF parameter. This value is defined in the SNA start list. These buffers hold data transmitted to and from SNA.	Used in questions 4 & 10			3
3	Number of channel-attached communication controllers (NCP) activated and owned by this SNA.	COMMON STORAGE: (NCPs * 1200)	= _____		5
4	Sum value of MAXBFRU parameters for all channel-attached communications controllers activated by this SNA. MAXBFRU is defined in the HOST definition statement of the NCP channel-attached major node definition.	COMMON STORAGE: ((IOBUF size + 98) * 2 * maxbfu NCPs)	= _____		3 & 6
5	Number of PUs defined in this SNA. Include all PUs defined to SNA in PU definition statements, and those controllers defined in CLUSTER definition statements. Include locally attached, remotely attached, dynamically added, switched, ICA, and NCP (including NTRI) PUs.	PRIVATE STORAGE: (defined PUs * 850)		= _____	48 + 67

Table 70. Worksheet for subarea communication management configuration (continued)					
Step	Description	Formulas (decimal)	Common	Private	DISPLAY STATS ID
6	Number of device type LUs defined in this SNA. Include the locally attached LUs, and the LUs attached through an NCP. In addition, include the LUs defined to SNA in LU definition statements, and those devices defined in TERMINAL statements. Do not include applications.	PRIVATE STORAGE: (defined LUs * 820)		= _____	50
7	Number of independent LUs either locally, remotely, or CDRSC defined. Include all independent LUs for which SNA provides boundary function services and all NTRI independent LUs.	COMMON STORAGE : (indep Lus * 270) PRIVATE STORAGE: (indep Lus * 400)	= _____	= _____	80
8	Number of device type LUs owned by this SNA in session with an application program owned by this SNA.	COMMON STORAGE: (LUs w/appls * 500) PRIVATE STORAGE: (LUs w/appls * 420)	= _____	= _____	71
9	Number of device type LUs owned by this SNA in session with an application program in another SNA.	PRIVATE STORAGE: (cross node LU ses * 400)		= _____	73
10	Number of device type (nonapplication) LUs in session with a TSO application.	COMMON STORAGE: (TSO LUs) * (2300 + IOBUF size)	= _____		55 & 3

Table 71. Summary of subarea communication management configuration		
CMC configuration description	Common	Private
Calculated COMMON storage for CMC configuration from above	= _____KB (<i>totcom</i>)	
DYNAMIC COMMON STORAGE (4 x <i>totcom</i>)	= _____KB	
SNA BASE STORAGE (COMMON)	+ 2953KB	
TOTAL SNA COMMON STORAGE	= _____KB	
Calculated PRIVATE storage for CMC configuration from above		= _____KB (<i>totpri</i>)
DYNAMIC PRIVATE STORAGE (2 x <i>totpri</i>)		= _____KB
SNA SYSPLEX STORAGE (if used)	+150KB	
SNA BASE STORAGE (PRIVATE)		+ 5507KB
TOTAL SNA PRIVATE STORAGE		= _____KB

APPL EAS storage estimates

EAS is the estimated number of concurrent sessions this application program will have with other logical units (LU-LU sessions). Accurate coding of the EAS value for your applications can save storage in your system. If your EAS value is specified as lower than the number of sessions that you actually have, sessions would still be established as usual. However, the efficiency of searching for the session representation could be impaired if a smaller table was allocated because of the lower EAS value.

For example, if you estimate that there will be less than 30 sessions with this application, but you let the EAS value default to 509, then an extra 4K table will be allocated from common storage. The size of the table is based on the EAS value that you code and is determined as follows:

Table size	EAS value
4K	30–4000
8K	4001–8000
16K	8001–16000
32K	16001–32000
64K	32001–48000
128K	48001–56000
256K	56001–64000
512K	greater than 64000

Note: For each OPEN ACB with EAS less than 30, one LF buffer will be allocated (for an LUCB) and there will be no additional table storage allocated. If EAS is greater than 30, then one SF buffer will be allocated (for an LUCB) in addition to storage for the table size indicated in the table above.

See [z/OS Communications Server: SNA Resource Definition Reference](#) for more information about EAS.

Appendix C. Communications storage manager

Communications storage manager (CSM) is a component of VTAM that allows authorized host applications to share data with VTAM and other CSM users without having to physically pass the data. A CSM user can be any system-authorized application program or product.

CSM is provided as part of the high-performance data transfer (HPDT) family of services. HPDT optimizes system performance for the transfer of bulk data. By providing a means for authorized applications to share buffers, CSM improves system performance during the transfer of bulk data by reducing the processing required for data movement. As a result, CPU resources (CPU cycles, memory bus, and cache) are conserved.

VTAM uses CSM to perform channel I/O over an HPDT MPC connection and to provide the HPDT services for host LU 6.2 applications. (HPDT MPC connections are explained in [“Multipath channel connections”](#) on page 40.) Host LU 6.2 applications use the IVTCSM and APPCCMD macroinstructions to reduce the use of system resources for large data transfers. For more information about HPDT services, see [“High-performance data transfer \(HPDT\)”](#) on page 339.

CSM includes an application programming interface (IVTCSM macroinstruction) that allows users to obtain and return storage in the form of CSM buffer pools by using the IVTCSM macroinstruction. Applications must be authorized to use the IVTCSM macroinstruction.

CSM buffer pools are identified by the fifteen distinct combinations of storage type and buffer size as described in [Table 72 on page 581](#) and specified on the IVTCSM REQUEST=CREATE_POOL macroinstruction.

Table 72. Buffer pools in CSM					
Storage types	Buffer sizes				
31-bit backed data space	4 KB	16 KB	32 KB	60 KB	180 KB
64-bit backed data space	4 KB	16 KB	32 KB	60 KB	180 KB
ECSA	4 KB	16 KB	32 KB	60 KB	180 KB

Data space storage is a common area data space and is associated with the master scheduler address space. This association results in a data space that persists for the life of the system.

Data space storage is either 31-bit backed or 64-bit backed. 31-bit backed data space, when fixed, resides below the 2 GB real storage bar. 64-bit backed data space, when fixed, can reside below or above the 2 GB bar. Where the storage is backed is a concern for only those products performing I/O into or out of the storage.

The application programming interface for CSM is described in [z/OS Communications Server: CSM Guide](#).

The rest of this appendix summarizes the interfaces that you can use to install, configure, monitor, and diagnose CSM storage.

CSM installation and definition

CSM is shipped and installed with the VTAM product tape. However, many CSM functions are independent of VTAM. CSM storage limits and tuning parameters are defined in the CSM parmlib member, IVTPRM00.

The first three lines in the CSM parmlib member define the maximum amount of storage to be dedicated to fixed, ECSA, and HVCOMM storage buffers in CSM. You can also specify one POOL definition for each CSM buffer pool of a particular *bufsize* and *bufsource* combination. On each POOL definition, you can specify:

- The minimum number of buffers to be free in the pool at any time (MINFREE). The storage pool is expanded when the number of free buffers falls below this limit.

- The number of buffers by which the pool is expanded (EXPBUF) when the number of free buffers falls below MINFREE.
- The initial number of buffers to be created in the pool (INITBUF). The pool does not contract below the contraction threshold value. CSM determines the contraction threshold value using the following formula:

$$\text{MAX}(\text{INITBUF}, (\text{MINFREE} + (2 * \text{EXPBUF})))$$

If parameters are not provided for a given CSM buffer pool, the IBM-supplied default values are used unless an application has provided these values on a IVTCSM REQUEST=CREATE_POOL macro.

See [z/OS Communications Server: New Function Summary](#) for a complete description of the CSM parmlib member.

IBM Health Checker for z/OS can be used to check whether appropriate values are defined for the maximum amount of storage to be dedicated to fixed and ECSA buffers in CSM. For more details about IBM Health Checker for z/OS, see [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#).

Initializing CSM

CSM is initialized by the first request to create a pool of buffers and remains active for the life of the system, independent of VTAM status. This could be issued by VTAM or a host application. Upon initialization, CSM reads the CSM parmlib member to determine storage limits and buffer pool related values.

Monitoring CSM

System operators can monitor the use of CSM storage by issuing the DISPLAY CSM command. CSM messages always start with the message prefix IVT. For a complete list of messages issued by CSM, see [z/OS Communications Server: SNA Messages](#). An application can monitor the use of CSM storage by using the SNA network monitoring NMI. See "SNA network monitoring NMI" in [z/OS Communications Server: IP Programmer's Guide and Reference](#) for more detailed information.

The following information is provided by the DISPLAY CSM command.

- Amount of storage allocated to each pool
- Amount of storage allocated to each user of the pool
- If OWNERID=ALL, the cumulative storage allocated to each user across all pools
- If the OWNERID value is not specified, the DISPLAY CSM command provides the following information:
 - The maximum amount of fixed, ECSA, and HVCOMM storage that can be allocated by CSM and current values for fixed, ECSA, and HVCOMM storage.
 - The highest level of fixed storage that is obtained since the last DISPLAY CSM command was issued without the OWNERID parameter.
 - The highest level of fixed storage that is obtained since the IPL.
 - The highest level of ECSA storage that is obtained since the last DISPLAY CSM command was issued without the OWNERID parameter.
 - The highest level of ECSA storage that is obtained since the IPL.
 - The highest level of HVCOMM storage that is obtained since the last DISPLAY CSM command was issued without the OWNERID parameter.
 - The highest level of HVCOMM storage that is obtained since the IPL.
 - The names of the CSM data spaces.

The DISPLAY CSM command can be used to identify a user of the pool that is consuming inordinate amounts of storage. This could occur in situations where an application fails to free buffers that it obtained from CSM. The report of storage allocated to a user is based on the user *owner_ID* (OWNERID

operand on the DISPLAY CSM command). CSM uses the application address space identifier (ASID) as the user OWNERID. The information by OWNERID indicates the amount of storage that must be freed by the user to enable the storage to be returned to the buffer pool.

CSM issues messages when CSM storage limits are either at a constrained level (approaching to 85% of defined limits) or at a critical level (90% of defined limits) or exceeded. In this case, the system operator can issue the MODIFY CSM command to increase the amount of fixed, ECSA, or HVCOMM storage that is available for CSM. You can issue MODIFY CSM command to increase or decrease the CSM storage limits without a re-IPL operation.

A Display CSMUSE command is available to provide the lower level of detail regarding CSM storage usage. The Display CSMUSE command provides displays of CSM storage usage by z/OS Communications Server components. Display of this information is intended to improve the serviceability of z/OS Communications Server and to aid in CSM storage diagnosis.

For more information about these CSM commands, see [z/OS Communications Server: SNA Operation](#). CSM storage information is also provided to the performance monitor interface (PMI). This information is equivalent to the information provided for the summary format of the DISPLAY CSM command. See [z/OS Communications Server: SNA Customization](#) for more information.

CSM problem diagnosis

You can obtain trace output of application requests to CSM by using the CSM option on the VTAM internal trace (VIT) or, when VTAM is not active, the GTF trace facility.

- When VTAM is operational, the CSM trace facility is controlled using the VIT. CSM writes records to the VIT using VTAM trace interfaces. The CSM trace option is used to control the generation of CSM trace records for both internal and external tracing.
- When VTAM is not operational, the VIT is not available and only external tracing is provided. The external trace is generated using the VTAM GTF event ID to write trace records directly to GTF in the same format as those recorded using VIT.

CSM tracing records the parameter list information that flows across the CSM interface and key internal events (such as pool expansion and contraction) for functions that manipulate buffer states. This allows you to trace and analyze the usage history of a buffer.

The number of trace records required to represent one IVTCSM request is variable based on the number of buffer operations requested. Because the information required to trace one IVTCSM request may span several CSM trace records, you can use the trace record flag field to determine whether additional trace records exist for a particular IVTCSM request. If the first bit in the trace record flag field is set on, the trace record is continued. If the VIT is not active, multiple trace records for an IVTCSM request could be interspersed with trace records of IVTCSM requests from other users. A unique trace record number is provided to correlate the continuation trace records for each IVTCSM request.

For more information about tracing events over the CSM API, see [z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures](#).

Appendix D. Logon manager

The transaction processing facility (TPF) is an application program dedicated to high-speed transaction processing. It is designed to handle a large volume of transactions submitted from many remote logical units. The logon manager is a VTAM application program that provides logon services for TPF.

The logon manager provides logon access to a TPF application program from independent and dependent logical units. If you define your TPF application program as a host processor (type 5), only dependent LUs can access the TPF application program. However, if you define your TPF application program as a type 2.1 physical unit, both dependent and independent LUs can access the TPF application program.

Note: To implement the logon manager, you need to have NCP Version 4 Release 3 or a subsequent release.

This appendix contains the following sections:

- How the Logon Manager Operates
- Installing the Logon Manager
- Starting the Logon Manager
- Defining the Logon Manager
- Monitoring Logon Manager Resources
- Halting the Logon Manager

How the logon manager operates

How the logon manager receives control depends on whether a dependent or independent LU is accessing the TPF application program. If the logon request is from a dependent SLU, the logon manager receives control when VTAM drives its logon exit. If the logon request is from an independent PLU, the logon manager receives control when VTAM drives the USERVAR exit, ISTEEXCUV. ISTEEXCUV is driven because the logon manager creates a USERVAR (with UVEXIT=YES) to represent the generic name of each TPF application program.

After the logon manager receives control, the control point LUs that support the generic application are identified. Each pair of supporting control point LUs is then passed to the CLU search exit, ELMCLUEX, to determine the LU best able to provide the service. The selection criteria of the ELMCLUEX exit can be determined by the user or the default can be used. The default criteria are:

- LU is below the session limit
- LU has the lowest hop count to the device
- LU has the fewest number of active sessions

For dependent LU access, the TPF application program can be defined as either a type 5 or type 2.1 PU.

For independent LU access, the TPF application program must be defined as a type 2.1 PU.

Installing the logon manager

The VTAM installation process places the logon manager load modules in SYS1.LINKLIB. You can put the logon manager in another load library by modifying the installation process. If you do this, you must either add this other load library to your system link list or point the logon manager JCL to this library by using a STEPLIB DD statement. Also, the logon manager must be APF authorized.

To run the logon manager, you must define the logon manager program ELMNGR as nonswappable with a protection key value of 6. To do this, update the MVS program properties table. See SCHEDxx macroinstruction in [z/OS MVS Initialization and Tuning Guide](#) and [z/OS MVS Initialization and Tuning Reference](#).

Starting the logon manager

To start the logon manager, issue the START command `START procname`, where *procname* is the name of the logon manager JCL procedure. The logon manager can also run as a job. If you want to run it as a job, modify the logon manager JCL and the start procedure to conform to your installation standards for such jobs.

You must place JCL for the logon manager in SYS1.PROCLIB or an equivalent library with a member name *procname* of your choice. Following is a sample of the required JCL:

```
//name      PROC MEMBER=partitioned data set member name,
              MAXSUBA=value
//GO        EXEC PGM=ELMNGR,PARM='MEMBER=&MEMBER,MAXSUBA=&MAXSUBA'
//ELMDEFDS  DD   DSN=partitioned dataset name,DISP=SHR
```

Note: The quotation marks are required around the PARM field of the EXEC statement.

name

Optional. If specified, it is usually the same as the member name *procname* of this JCL procedure.

ELMNGR

The name of the logon manager load module.

ELMDEFDS

The required DD name of the DD statement for the partitioned data set that contains the logon manager configuration definition.

MEMBER

The 1 to 8 character name of the member of the ELMDEFDS partitioned data set that contains the logon manager configuration definition.

MAXSUBA

Specifies the maximum number of network subareas that can be defined to the logon manager. Valid values are in the range 1 – 511. The default is 1. The number specified must be at least as great as the number of network subareas defined in the logon manager configuration definition.

Defining the logon manager

This section describes how to define the logon manager and includes a description of a typical logon manager configuration.

Note: You can also define TPF as a type 5 data host, rather than a type 2.1 physical unit. As a data host, TPF supports only fixed pacing window sizes, virtual route 0 (VR=0), and transmission priority 2 (TP=2). The TPF data host must reside in its own SNI network with VTAM in another host acting as the gateway VTAM. In this type of configuration, you need to code the TPF application programs as dependent logical units. These are the only resources the TPF data host can own. When TPF is a type 5 data host, logon manager should not be used.

Sample logon manager configuration

Figure 151 on page 587 shows a logon manager configuration consisting of:

- TPF with control point LUs
- TPF defined as a type 2.1 PU
- VTAM with the logon manager
- Dependent SLUs

TPF System
(Type 2.1 Peripheral Node)

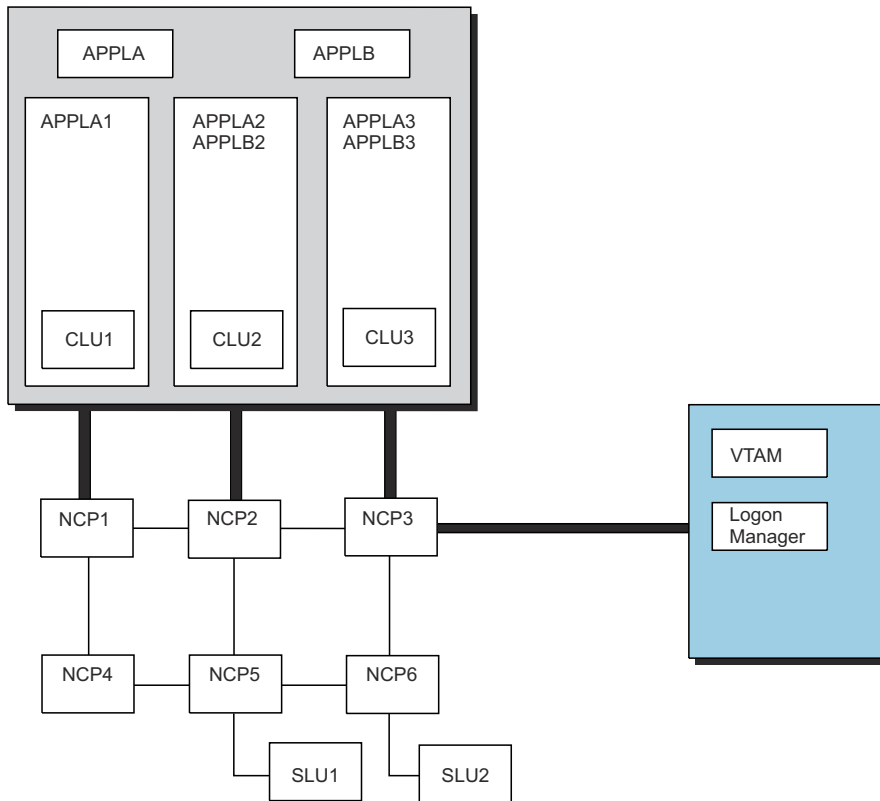


Figure 151. Sample Logon Manager Configuration

NCP1 through NCP3 are NCP Version 4 Release 3 subarea nodes, which are channel-attached to TPF. NCP3 is also channel-attached to VTAM. NCP1 through NCP6 are link-attached to each other. NCP4 through NCP6 do not have to be NCP Version 4 Release 3 or subsequent releases.

CLU1 through CLU3 are control logical units (CLUs). They manage the flow of transactions from their associated subarea nodes.

TPF is defined to VTAM as a type 2.1 peripheral node. Because the NCPs are channel-attached, TPF requires LNCTL=CA on the NCP GROUP definition statement. Interactions between logon manager and TPF occur in the form of control point LU to control point LU sessions.

APPLA and APPLB are TPF application programs. They are defined to VTAM as local applications.

Resources defined as independent LUs can initiate sessions with the TPF applications.

To take advantage of multiple-channel attachments into the subarea network, each application program (for example, APPLA) needs at least one LU per channel attachment. In this configuration, each application program requires three or more LUs (for example, APPLA1, APPLA2, APPLA3) to use all channels into the network. Because APPLB has only two LU definitions, it can use only two channels into the network.

SLU1 and SLU2 are defined to VTAM as dependent SLUs. They are link-connected to their associated subarea nodes and represent the remote terminal users of the TPF application programs.

You can define class of service (COS) tables for a subarea configuration. For information about COS tables, see [z/OS Communications Server: SNA Resource Definition Reference](#).

Defining the logon manager and TPF applications to VTAM

You need to define the logon manager and each TPF application program to VTAM using APPL definition statements. The following operands and operand values are required:

```
name      APPL  ACBNAME=acbname,
                AUTH=(ACQ,NOCNM,NOPASS,NOTSO,SPO),
                PARSESS=YES,
                SRBEXIT=NO,
                SONSCIP=YES,
                ENCR=NONE,
                VTAMFRR=NO,
                HAVAIL=YES
```

If *acbname* differs from the name *procname* of the logon manager JCL, you need to code the LMPROC definition statement in the logon manager configuration definition.

The following APPL definition statement operands are optional:

- PRTCT
- USSTAB
- VPACING

To use password protection for the logon manager, include the PRTCT operand on the APPL definition statement for the logon manager, and specify the same password as that on the LMPROC definition statement in the logon manager configuration definition. To use password protection for a TPF application program, include the PRTCT operand on the APPL definition statement for the TPF application program, and specify the same password as that on the LMAPPL definition statement in the logon manager configuration definition.

Defining the logon manager configuration

The logon manager configuration is defined by a sequence of source statements stored in a member of a partitioned data set with the member name that is specified in the logon manager JCL. You can write different types of source statements in any order in the configuration definition. Additional logon manager configuration definitions can be stored in members with other member names.

A logon manager configuration definition always contains one subarea (SA) definition statement identifying each network subarea. Other logon manager configuration definition statements are optional.

Subarea configuration

The system programmer responsible for identifying the network subareas and their adjacent subareas must provide a configuration of the network subareas. Each network subarea must be defined in the logon manager configuration definition using one of the following formats:



x
The address for a subarea (SA).

a
The address for an adjacent subarea (ADJSA). The subarea must be adjacent to subarea x.

You can specify multiple adjacent subarea addresses by enclosing them within parentheses and separating them with a comma.

All network subareas that you are defining to the logon manager must have their addresses appear in SA definition statements as either *x* or *a*.


Note: The value of MAXSUBA specified on the EXEC statement in the logon manager JCL must be at least as great as the number of different network subareas addressed as x or a in the SA definition statements in the logon manager configuration definition.

The SAs supplied by the system programmer are used by the logon manager to build a HOP table. The HOP table contains the minimum distance between subareas and is used for multitask support.

Class of service configuration

With logon manager, you can use class of service (COS) to distinguish between the needs of various sessions. Logon manager differentiates between session types based on the COS name. By designating different COS names for batch and interactive traffic (or any other preferred distinction), you can keep the various types of traffic from interfering with each other. Using this function, you define multiple views of the network that list preferred paths for different types of session traffic.

To specify the COS name, use the COS definition statement in the following format:



For example, in Figure 151 on page 587, NCP6 is only 1 hop away from NCP3, but is 2 hops away from NCP2. You can then specify that NCP3 is the preferred boundary function for any interactive sessions to the host and that NCP2 is the preferred route for batch traffic. To define the different routes to logon manager, you can define a number of Classes of Service, each with its view of the network. In the previous example you would code one Class of Service for interactive traffic, and one for batch.

To define these different views, use the COS definition statement. For example, to code the interactive view described previously, you would code the following statements:

```
COS=INTERACT
SA=6, ADJSA=(3)
```

Using this technique, NCP3 would be used for interactive traffic to the host.

The result is that each session is assigned to a TPF server based on predefined routing requirements, and enables you to assign suitable routes.

The SA definition statements that follow a COS definition statement apply only to the table representing that COS. If you code an SA definition statement without a COS definition statement preceding it, that SA definition statement is used as part of the default table.

If a COS name is not found for a COS that is requested, logon manager uses the default tables for the session:

- If no COS name is specified for the session, the default table is used.
- If a COS name is for the session, but cannot be found, the default table is used.
- If the default table is to be used, but is not defined to logon manager, an error is reported.

In the following example, coded for the configuration in Figure 151 on page 587, the first entry is used as part of the default table, and the other entries are used for interactive and batch-directed traffic.

```
SA=1, ADJSA=(2)
COS=INTERACT
SA=6, ADJSA=(3)
SA=5, ADJSA=(2)
COS=BATCH
SA=6, ADJSA=(2)
SA=5, ADJSA=(3)
```

Application passwords and minimum channel links

To define password protection for a TPF application program, include the following definition statement in the logon manager configuration definition. You can also use this definition statement to define the number of active channel links required before the application program is supported by the logon manager.

▶ LMAPPL — = — application program name

, — PASSWD — = — password

. — MINLINK — = — value

application program name

The name of the TPF application program. It must be the name of a real application program.

password

The 1 to 8 character password for the TPF application program. Passwords are optional.

If password protection is required for a TPF application program, the password must be specified on the PASSWD operand on the LMAPPL definition statement and on the PRTCT operand on the VTAM APPL definition statement. There is no default password.

value

Defines the minimum number of channel links that must be active before the application program can engage in a session. Valid values are int in the range 1 – 255. The default is 1. This definition is optional.

If all application programs on the TPF have the same password and the same required number of active channel links, you can use the following definition statement:

LMAPPL — = — **ALL** — , — **PASSWD** — = — *password*

. — **MINLINK** — = — *value*

Logon manager procedure

If the logon manager requires password protection or if the logon manager ACBNAME on the APPL definition statement differs from the MVS *procname* or *jobname*, include the following definition statement in the logon manager configuration definition:

LMPROC — = — *acbname* ————— **PASSWD** — = — *password*

acbname

The 1 to 8 character name of the logon manager's ACB. If you do not supply the LMPROC definition statement, the MVS *procname* or *jobname* is used.

password

The 1 to 8 character password for the logon manager. The logon manager uses this password when opening the ACB for the *procname* or *jobname*. The logon manager password is optional.

If password protection is required for the logon manager, the password must be specified on the PASSWD operand on the LMPROC definition statement and on the PRTCT operand on the VTAM APPL definition statement. There is no default password.

Number of TPF applications

To specify the number of TPF application programs that can be defined, include the following definition statement in the logon manager configuration definition:

➤➤ MAXAPLC — = — *value* ➤➤

value

The number of TPF application programs that can be defined. Valid values are in the range 1 – 4095. The default is 4. This definition is optional.

Maximum number of buffers

To specify the maximum number of buffers to allocate for each channel-attached control point LU, include the following definition statement in the logon manager configuration definition:

➤➤ MAXBUF — = — *value* ➤➤

value

The maximum number of buffers to allocate for each control point LU. Valid values are in the range 2 – 65 535. The default is 30. This definition is optional.

Logon manager takes the value specified in the MAXBUF definition statement and multiplies it by the value specified in the CHANATT definition statement to determine the total number of buffers available to logon manager. There is no restriction to the number of buffers that can be used for a given control point LU as long as there are free buffers available.

Maximum number of subareas

To specify the maximum number of channel-attached subareas, include the following definition statement in the logon manager configuration definition:

➤➤ CHANATT — = — *value* ➤➤

value

The maximum number of channel-attached subareas. This number is used to determine the maximum number of channel-attached control point LUs. Valid values are in the range 1 – 255. The default is 4. This definition is optional.

Reserved keywords

The following keywords are reserved for use with the TPF logon manager:

- ADJSA
- ALL
- APPL
- CHANATT
- CLU
- HELP
- ID
- INFO
- LMAPPL
- LMPROC
- MAXAPLC
- MAXSUBA

- MEMBER
- MINLINK
- PASSWD
- PEND
- SA
- STOP

Monitoring logon manager resources

The VTAM logon manager application offers commands to monitor the status of the logon manager and the resources under control of the logon manager.

By using the logon manager MODIFY command, you can display information about logon manager resources (TPF applications, CLUs, and LUs), including session status. For more information about the logon manager operator commands, see [z/OS Communications Server: SNA Operation](#).

Halting the logon manager

The logon manager functions as a stand-alone VTAM application and, thus, can be stopped independently of VTAM. The logon manager MODIFY STOP command stops logon manager processing. For more information about the logon manager operator commands, see [z/OS Communications Server: SNA Operation](#). When the command is issued, the logon manager shuts down its subtasks, closes its ACBs, and returns control to MVS. Current sessions are allowed to end, but new sessions are not permitted.

The VTAM HALT, HALT QUICK, and CANCEL commands terminate the logon manager, as they do any VTAM application.

Appendix E. Cryptographic keys

If you use the VTAM data encryption facility, you need to file cryptographic keys on the cryptographic key data set at the appropriate host processors. For information about which hosts require cryptographic keys, see [“Cryptography facility” on page 289](#).

This appendix describes how to file these keys for different types of cryptographic facilities for both single-domain and multiple-domain sessions. The available cryptographic services are:

- z/OS Integrated Cryptographic Service Facility (ICSF) and S/390® or zSeries Cryptographic Co-Processor
ICSF is a licensed program that runs under MVS and provides access to the hardware cryptographic feature for programming applications. The combination of the hardware cryptographic feature and ICSF provides secure high-speed cryptographic services.
- Other PCF/CUSP or Common Cryptographic Architecture (CAA) compatible cryptographic products

Note: Triple-DES 24-byte encryption requires the use of the ENCRYPTN=CCA start option and that the Common Cryptographic Architecture (CCA) product is present. Otherwise, sessions that require triple-DES 24-byte encryption will fail. CCA defines a set of cryptographic functions, external interfaces, and a set of key management rules that provide a consistent, end-to-end cryptographic architecture across different IBM platforms.

The following references are used with compatible cryptographic products:

PCF/CUSP

Refers to any cryptographic product that is compatible with PCF/CUSP.

CCA

Refers to any cryptographic product that is compatible with Common Cryptographic Architecture (CCA).

Notes:

1. If ICSF/MVS runs in CUSP mode, use the information for PCF/CUSP.
2. When using ICSF in PCF compatibility mode and migrating from an existing PCF cryptographic key data set (CKDS), an importer key with a key value of the PCF master key value must be included. Use the PCF master key 8 bytes twice to create the ICSF 16-byte key. See the ICSF publications for additional information.

Specific commands and control statements for key input may differ by product.

For more information about establishing cryptographic sessions, see [z/OS Communications Server: SNA Programming](#).

Filing SLU keys for single-domain cryptographic sessions

The process for filing SLU keys for single-domain cryptographic sessions varies, depending on which cryptographic service you use.

Single-domain cryptographic sessions that use PCF/CUSP

To allow cryptographic sessions that use PCF/CUSP to be established within a single domain, install the PCF/CUSP compatible cryptographic product.

Use the PCF/CUSP compatible cryptographic product to file secondary logical unit (SLU) keys on the cryptographic key data set (CKDS) as follows:

- For each device-type LU that is to be used as the secondary end of a cryptographic session, code the following statement:

```
LOCAL luname
```

where *luname* is the name of the LU.

This LOCAL statement generates an SLU key for the LU and adds it to the CKDS enciphered under the first variant of the host master key. It also returns a clear SLU key. Enter the clear SLU key into the device.

The CKEYNAME operand on the LU definition statement can be used to reduce definition statements coded in the CDKS. Multiple devices can use the same LOCAL statement in the CDKS by specifying the same value for the CKEYNAME operand.

```
LOCAL ckeyname
```

where the value specified in the CKEYNAME operand is the same as the *ckeyname*.

See the [z/OS Communications Server: SNA Resource Definition Reference](#) for information about coding the CKEYNAME definition statement.

- For each VTAM application program that is to be the secondary end of a cryptographic session, code the following statement:

```
REMOTE name
```

where *name* is the name of the application program.

This REMOTE statement generates an SLU key for the application program and adds it to the CKDS enciphered under the second variant of the host master key.

Single-domain cryptographic sessions that use ICSF/MVS

To allow cryptographic sessions that use ICSF/MVS to be established within a single domain, install ICSF/MVS.

Use ICSF/MVS to file SLU keys on the CKDS as follows:

- For each device-type LU that is to be used as the secondary end of a cryptographic session, code the following, where *name* is the name of the LU:

```
ADD LABEL(name) TYPE(EXPORTER) CLEAR SINGLE
```

Note: SINGLE is needed only if the cryptographic support of the device-type LU does not support double length keys.

This key generation utility program (KGUP) statement generates an exportable SLU key for the LU and adds it to the CKDS that is enciphered under the host master key. It also returns a clear SLU key. Enter the clear SLU key into the device.

- For each VTAM application program that is to be used as the secondary end of a cryptographic session, code the following, where *name* is the name of the application program:

```
ADD LABEL(name) TYPE(IMPORTER) CLEAR SINGLE
```

Note: SINGLE is needed only if the cryptographic support of the device-type LU does not support double length keys.

This KGUP statement generates an importable SLU key for the application program and adds it to the CKDS that is enciphered under the host master key.

Filing CDRM keys for cross-domain cryptographic sessions

The process for filing SLU keys for cross-domain cryptographic sessions varies, depending on which cryptographic service you use and whether both hosts use the same cryptographic service.

Note: CCA processing must be present to use triple-DES encryption. The ENCRYPTN=CCA start option is also required for triple-DES encryption. PCF/CUSP compatible cryptographic products typically support DES but do not support triple-DES.

Cross-domain cryptographic sessions in which both hosts use PCF/CUSP

To allow cross-domain cryptographic sessions to be established, file SLU keys for each domain as described in “Single-domain cryptographic sessions that use PCF/CUSP” on page 593. Start the PCF/CUSP compatible cryptographic product before you activate the external CDRM for which cross statements have been filed in the cryptographic keys data set. Use the PCF/CUSP compatible cryptographic product to file cross-domain keys on the cryptographic key data set (CKDS) at each host processor as follows:

- For each pair of host processors (HOST1 and HOST2) that are to have cross-domain cryptographic sessions between their domains, code the following at HOST1, where *name* is the name of HOST2 CDRM:

```
CROSS name
```

This CROSS statement generates two cross-domain keys, one defined as local and the other defined as remote.⁴ It stores the local cross-domain key in the CKDS enciphered under the first variant of HOST1 host master key and stores the remote cross-domain key in the CKDS enciphered under the second variant of HOST1 host master key. Both of these keys are associated with the name of HOST2 CDRM. This CROSS statement also returns clear copies of the two keys.

- The cross-domain keys generated at HOST1 must be used at HOST2 and supplied as input to PCF/CUSP in a CROSS statement.

When HOST2 has PCF

```
CROSS name, KEYLOC=x, KEYREM=y, ADD
```

When HOST2 has CUSP

```
CROSS name, KEYLOC=x, IKEYLOC=xx, KEYREM=y, IKEYREM=yy, ADD
```

where:

name is the name of HOST1 CDRM

x is the clear remote cross-domain key from HOST1

xx is the second half of double key value from HOST1 remote key

y is the clear local cross-domain key from HOST1

yy is the second half of double key value from HOST1 local key

This CROSS statement does the following actions:

- Adds the two cross-domain keys to HOST2 CKDS
- Associates the two keys with the name of HOST1 CDRM
- Reverses the local/remote relationship of the two keys
- Enciphers key *x* and its intermediate key *xx* under the first variant of HOST2 host master key, and key *y* and its intermediate key *yy* under the second variant of HOST2 host master key

Figure 152 on page 596 illustrates the possible coding for cross-domain cryptographic sessions where both hosts use PCF/CUSP. APPL2A can initiate a cryptographic session with LU1A. APPL2A can be the primary LU (PLU) in the cryptographic session with APPL1A. However, APPL2A cannot be the secondary LU in a cryptographic session with APPL1A because there is no REMOTE statement for APPL2A in VTAM2.

⁴ The terms *local* and *remote* in reference to the keys used by the CROSS statement do not have the same meaning here as they do in other contexts in this document. For more information about PCF, see the [z/OS Cryptographic Services ICSF Application Programmer's Guide](#). For more information about CUSP, see the [z/OS Cryptographic Services ICSF Administrator's Guide](#).




Figure 152. Cryptography in multiple-domain environment (Both hosts use PCF/CUSP)

For the configuration in Figure 152 on page 596 to have any encrypted sessions, start cryptography in both hosts before activating a session.

Cross-domain cryptographic sessions in which both hosts use ICSF/MVS

To allow cross-domain cryptographic sessions to be established, file SLU keys for each domain as described in “Single-domain cryptographic sessions that use ICSF/MVS” on page 594. Then use ICSF/MVS to file cross-domain keys on the cryptographic key data set (CKDS) at each host processor as described in the following paragraphs.

A complementary pair of exporter and importer keys must be generated for the two host processors (HOST1 and HOST2). To allow for cross-domain cryptographic sessions between hosts using the ICSF/MVS cryptographic service, perform the following steps:

- At HOST1, code the following key generation utility program (KGUP) statements:

```
ADD LABEL(name) TYPE(EXPORTER) CLEAR
ADD LABEL(name) TYPE(IMPORTER) CLEAR
```

where *name* is the name of HOST2 CDRM

These KGUP statements generate an exporter key for HOST1 to use to encrypt session keys sent from HOST1 to HOST2 and an importer key for HOST1 to decrypt session keys sent from HOST2 to HOST1. The exporter and importer keys are placed in HOST1 CKDS. The statements also cause the clear key values to be placed in the KGUP key output data set. Assume clear key value X,XX for the exporter key and clear key value Y,YY for the importer key. These values are to be used on HOST2 to build complement key control statements for these keys.

- At HOST2, code the following KGUP statements using the clear key values from HOST1:

```
ADD LABEL(name1) TYPE(IMPORTER) CLEAR KEY(x,xx)
ADD LABEL(name1) TYPE(EXPORTER) CLEAR KEY(y,yy)
```

where:

name1

is the name of HOST1 CDRM.

x

is the first half of the double-length clear key value of HOST2 exporter key generated in HOST1.

xx

is the second half of the double-length clear key value from HOST2 exporter key generated in HOST1.

y

is the first half of the double-length clear key value from HOST2 importer key generated in HOST1.

yy

is the second half of the double-length clear key value from HOST2 importer key generated in HOST1.

These control statements place these keys in HOST2 CKDS.

Figure 153 on page 597 illustrates the possible coding for cross-domain cryptographic sessions where both hosts use ICSF/MVS. APPL2A can initiate a cryptographic session with LU1A. APPL1A can initiate a cryptographic session with LU2A. APPL2A can be the PLU in a cryptographic session with APPL1A. However, APPL2A cannot be the SLU in a session with APPL1A because there is no TYPE(IMPORTER) statement coded in VTAM2, and APPL2A requires cryptography when it is the SLU.

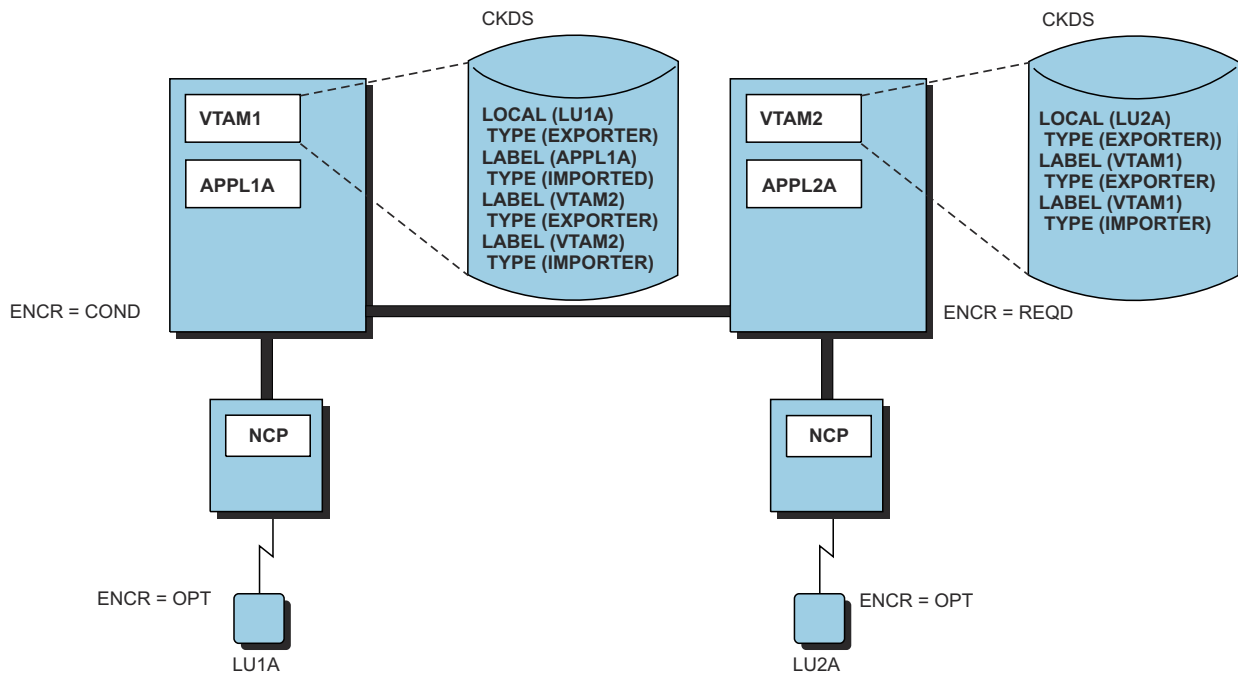


Figure 153. Cryptography in multiple-domain environment (Both hosts use ICSF/MVS)

For the configuration in Figure 153 on page 597 to have any encrypted sessions, start cryptography in both hosts before activating a session.

For more information about ICSF/MVS, see the [z/OS Cryptographic Services ICSF Administrator's Guide](#).

Cross-domain cryptographic sessions in which the hosts use different cryptographic services

In this description of possible coding, assume that ICSF/MVS is installed in HOST1 and a PCF/CUSP compatible cryptographic product is installed in HOST2. A complementary pair of key-encrypting keys must be generated for HOST1 and HOST2. To allow for cross-domain cryptographic sessions between the hosts, perform one of the following, depending upon where key processing is to be performed first:

- When key processing is to be performed from HOST2 first, perform these steps:
 1. At HOST2 where the PCF/CUSP compatible cryptographic product is installed, code the following key generation utility program (KGUP) statement:

```
CROSS name
```

Where *name* is the name of HOST1 CDRM.

This PCF/CUSP KGUP statement generates a LOCAL and REMOTE key pair and places them in HOST2 cryptographic key data set (CKDS). A copy of the key pair generated will also appear in the KGUP SYSPRINT output listing. This gives HOST2 the capability to send encrypted session keys to HOST1.

2. At HOST1, use the PCF/CUSP compatible cryptographic product SYSPRINT output from the cross key pair to create the following ICSF control statement for the appropriate environment.

When HOST2 has PCF

```
ADD LABEL(name) TYPE(IMPORTER) CLEAR NOCV KEY(key-value)
```

When HOST2 has CUSP

```
ADD LABEL(name) TYPE(IMPORTER) CLEAR NOCV KEY(key-value,ikey)
```

where:

name is the name of HOST2 CDRM.

key-value is the key value from HOST2 of the local key.

ikey is the intermediate key value from HOST2 of the intermediate local key.

This ICSF KGUP statement places the importer key, which is the same clear value as HOST2 LOCAL key in HOST1 CKDS. This allows HOST1 to decrypt session keys sent from HOST2.

3. At HOST1, use the PCF/CUSP compatible cryptographic product SYSPRINT output from the cross-key pair to create the following ICSF control statement for the appropriate environment.

When HOST2 has PCF

```
ADD LABEL(name) TYPE(EXPORTER) CLEAR NOCV KEY(key-value)
```

When HOST2 has CUSP

```
ADD LABEL(name) TYPE(EXPORTER) CLEAR NOCV KEY(key-value,ikey)
```

where:

name is the name of HOST2 CDRM.

key-value is the key value from HOST2 of the remote key.

ikey is the intermediate key value from HOST2 of the intermediate remote key.

This ICSF KGUP statement places the exporter key, which has the same clear value as HOST2 REMOTE key, in HOST1 CKDS. This allows HOST1 to send encrypted session keys to HOST2.

- When key processing is to be performed from HOST1 first, perform these steps:

1. At HOST1 where ICSF/MVS is installed, code the appropriate key generation utility program (KGUP) statements listed below for the environment in HOST2:

When HOST2 has PCF

```
ADD LABEL(name) TYPE(EXPORTER) CLEAR NOCV SINGLE  
ADD LABEL(name) TYPE(IMPORTER) CLEAR NOCV SINGLE
```

When HOST2 has CUSP

```
ADD LABEL(name) TYPE(EXPORTER) CLEAR NOCV  
ADD LABEL(name) TYPE(IMPORTER) CLEAR NOCV
```

where *name* is the name of HOST2 CDRM.

These statements cause two keys to be generated in HOST1 CKDS. One key, the exporter, is to be used when sending session keys to HOST2. The other key, the importer, is to be used to decrypt session keys sent from HOST2.

The complement clear key values of both the exporter and importer are placed in the ICSF key output data set. Use these values as the clear key values for the CROSS statements on HOST2.

- At HOST2 where the PCF/CUSP compatible cryptographic product is installed, code the appropriate key generation utility program (KGUP) statement listed below for the environment in HOST2.

When HOST2 has PCF

```
CROSS name,KEYLOC=x,KEYREM=y,ADD
```

When HOST2 has CUSP

```
CROSS name,KEYLOC=x,IKEYLOC=xx,KEYREM=y,IKEYREM=yy,ADD
```

where:

name is the name of HOST1 CDRM.

x is the key value from HOST1 of the IMPORTER key.

y is the key value from HOST1 of the EXPORTER key.

xx is the second half of double key value from HOST1 IMPORTER key.

yy is the second half of double key value from HOST1 EXPORTER key.

Figure 154 on page 599 illustrates the possible coding for cross-domain cryptographic sessions where HOST1 uses ICSF/MVS and HOST2 uses PCF/CUSP. APPL2A can initiate a cryptographic session with LU1A. APPL1A can initiate a cryptographic session with LU2A. APPL2A can initiate a cryptographic session with APPL1A. However, APPL1A cannot be a PLU in a session with APPL2A because there is no TYPE(IMPORTER) statement coded in VTAM2, and APPL2A requires cryptography.

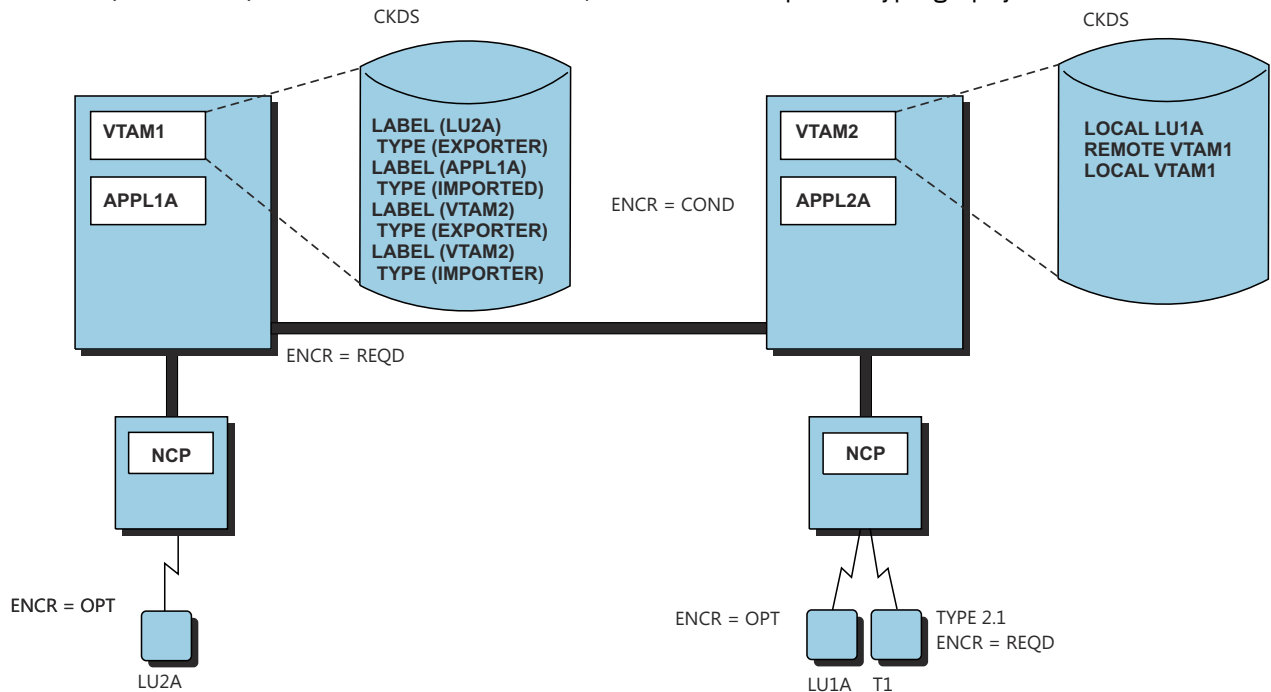


Figure 154. Cryptography in multiple-domain environment (Hosts use different cryptographic services)

For the configuration in Figure 154 on page 599 to have any encrypted sessions, start cryptography in both hosts before activating a session.

Changing the cryptographic capability of a logical unit

If the VTAM encryption facility is installed, the cryptographic capability of the logical unit can be changed for future sessions by using the MODIFY SECURITY or MODIFY ENCR command. The ENCRTYPE keyword can only be changed using the MODIFY SECURITY command. For the syntax of the command, see [z/OS Communications Server: SNA Operation](#).

In addition, an alternate LU key-encrypting key can be used while the master key is being changed in the CKDS. See [“Changing cryptographic keys dynamically” on page 295](#) for additional information.

Appendix F. Command lists: Dynamic configuration of channel-attached devices

The IBM-supplied command lists are used to illustrate dynamic configuration of channel-attached devices (dynamic I/O). If you want to modify the command lists, this section describes how to call the command lists and what output the command lists return to assist you during modification. For general information, see [“Dynamic configuration of channel-attached devices” on page 170](#).

ISTDINFO: VTAM device information services

The ISTDINFO command list accepts requests from a NetView REXX command list and uses the information to retrieve information about VTAM devices. ISTDINFO then returns the device information to the caller. You can call the ISTDINFO command list using following format:

```
ISTDINFO  DEV_INFO, dev_string
```

The comma separating the parameters is mandatory. The parameters provide the following information:

DEV_INFO

The character string "DEV_INFO", in either uppercase or lowercase.

dev_string

A character-string constant representing a hexadecimal device address. Device addresses consist of 1–4 hexadecimal characters (for example, X'F12').

Dependencies and restrictions

If you are using the ISTDINFO command list, you should be aware of the following dependencies and restrictions:

- You can only IPL the configuration using an HCD I/O Definition (IODF). You cannot use MVS/CP to define your configuration. The LOADxx parmlib member, which contains the IPL parameters, must have an IODF statement coded, not an MVS/CP statement.
- You can define only the following devices:
 - Local SNA clusters
 - Communication controllers
- The devices must be defined to MVS using the IBM-supplied unit information modules.
- Complete device information is returned only when the hardware is self-describing.

Output variable

The ISTDINFO command list returns information in the REXX special variable RESULT. RESULT is a string containing the following tokens:

```
return_status  
device_address  
device_class  
self_describing  
device_type  
device_model  
defined_type  
generic_family  
PU_type
```

TG_number
IPL_online
current_online
allocate
config_token
maxbufnu
maxdata
logical_terminals
owner
ncp_pu_type
ncp_timeout
ncp_delay
ncp_maxbufnu
ncp_inbfrs
ncp_maxdata
ncp_transfr
ncp_bfrpad
ncp_TPS_feature

Output tokens

The following describes the information contained in the tokens:

return_status

Returned status from the device information services routine:

- 0**
Device specific information returned
- 2**
Not valid option string specified for parameter one
- 3**
Not valid device address specified for parameter two
- 4**
No VTAM supported devices
- 5**
Unsupported operating system environment
- 101**
Internal error 101
- 102**
Internal error 102
- 201**
Internal error 201
- 301**
Internal error 301
- 302**
Internal error 302
- 303**
Internal error 303
- 304**
Internal error 304
- 305**
Internal error 305

306
Internal error 306

307
Internal error 307

308
Internal error 308

309
Internal error 309

310
Internal error 310

311
Internal error 311

401
Internal error 401

501
Internal error 501

502
Internal error 502

503
Internal error 503

504
Internal error 504

505
Internal error 505

601
Internal error 601

602
Internal error 602

603
Internal error 603

604
Internal error 604

605
Internal error 605

606
internal error 606

607
Internal error 607

608
Internal error 608

609
Internal error 609

610
Internal error 610

611
Internal error 611

device_address

Address of the device associated with the returned information. This is not necessarily the same as the input *dev_string*.

device_class

Device class stored in the unit control block (for example, COMM, DISP or CTC).

self_describing

An indication of whether the device is self-describing (YES or NO).

device_type

Actual device type returned to MVS during VARY ONLINE processing, or *N/A* if not applicable.

device_model

Model number for this device, or *N/A* if not applicable.

defined_type

Defined device type specified by the user. For example, 3705 is specified for a 3705, 3725, 3745, or *N/A* if not applicable. This should always be defined.

generic_family

Generic family of this device specified by the device: 3745 for 3745, or *N/A* if not applicable.

PU_type

Physical unit type of the device (for example, 2 or 21 or 4, or *N/A* if not applicable).

TG_number

Transmission group number of the device: 0–255, or *N/A* if not applicable.

IPL_online

Status of the device at IPL: ONLINE, OFFLINE, or *N/A* if not applicable.

current_online

Status of the device currently: ONLINE or OFFLINE.

allocate

Status of allocation: ALLOCATED or UNALLOCATED.

config_token

Token representing the current I/O configuration. The actual I/O configuration is not printable in EBCDIC.

maxbufru

Maximum number of buffers to be set aside by the host, or *N/A* if not applicable.

maxdata

Maximum number of bytes the PU can receive in one PIU, or *N/A* if not applicable.

logical_terminals

Character string of 256 ones and zeros. A 1 indicates a logical terminal exists at that offset, or *N/A* if not applicable. This string defines the number of logical units to be generated and the local address for each.

owner

Value entered into HCD for the subsystem or access method using this device, or *N/A* if not applicable. Any embedded blanks in the owner string are replaced with the underscore (_) character.

ncp_pu_type

Value entered into HCD for PU type, or *N/A* if not applicable.

ncp_timeout

Value entered into HCD for attention response timeout for the device, or *N/A* if not applicable.

ncp_delay

Value entered into HCD for maximum delay before attention to channel, or *N/A* if not applicable.

ncp_maxbufru

Value entered into HCD for maximum number of host buffers allocated for the device, or *N/A* if not applicable.

ncp_inbfrs

Value entered into HCD for initial allocation of control buffers for the device, or *N/A* if not applicable.

nbp_maxdata

Value entered into HCD for the maximum number of bytes the PU can receive in one data transfer for the device, or *N/A* if not applicable.

nbp_transfr

Value entered into HCD for the number of NBP buffers for MAXDATA transfer for the device, or *N/A* if not applicable.

nbp_bfrpad

Value entered into HCD for number of pad characters to the access method for the device, or *N/A* if not applicable.

nbp_TPS_feature

Value entered into HCD for mode of the two-processor switch feature for the device: TPS, TCS, NONE, or *N/A* if not applicable.

ISTDEFIN: VTAM device information services

The ISTDEFIN command list accepts one or more device addresses, gets information about each device, defines it to VTAM, ensures that each device is online, and activates the VTAM resource definitions. You can call the ISTDEFIN command list using the following format:

```
Call ISTDEFIN func_code dev_str
```

The parameters provide the following information:

func_code

One of the following function codes that indicate how the input device addresses correlate to the device processed by the command list:

1

Process only device addresses passed to this command list.

2

Use input device address as a base address to locate communication, display or CTC device to process.

dev_str

A string of one or more hexadecimal device addresses separated by blanks.

For example, if you wanted to call the ISTDEFIN command list with a function code of 1 and 3 device addresses, X'00A', X'123', and X'FE3', you would use the following command:

```
ISTDEFIN 1 A 123 FE3
```

Dependencies and restrictions

If you are using the ISTDEFIN command list, you should be aware of the following dependencies and restrictions:

- You need NetView 1.3 or later releases or a product providing a similar interface and services.
- ISTDEFIN must run in the NetView address space.
- You can define only the following devices:
 - Local SNA clusters
 - Communications controllers
- You need to concatenate USER1.AUTO.VTAMLST with SYS1.VTAMLST.
- You can define NCP major nodes only if they are channel-attached.
- The IBM-supplied naming convention supports only 36 subarea addresses.
- Error messages for the VARY NET,ACT command are diagnosed only to the extent that syntax errors in the resource definitions generated by this sample are identified.

Output variable

The ISTDEFIN command list returns information in the REXX special variable RESULT. RESULT is a string containing the following tokens:

```
return_status  
device_address1  
device_rc1  
device_address2  
device_rc2  
:  
device_addressn  
device_rcn
```

where *n* is the number of devices processed by the command list.

Output tokens

The following describes the information contained in the tokens:

return_status

Highest return code encountered while processing the device string.

device_address₁

Actual device address processed.

device_rc₂

Return code that resulted from processing device_address_i. The possible values are:

- 0**
device successfully defined and activated
- 1**
device address not found
- 2**
not valid option string (parameter 1) specified on ISTDINFO call
- 3**
not valid device address (parameter 2) specified on ISTDINFO call
- 4**
device is currently allocated
- 5**
operating environment not supported
- 6**
operator canceled processing for device
- 7**
timer expired while varying device online
- 8**
error occurred while updating USER1.AUTO.VTAMLST
- 9**
VARY ONLINE command failed for device
- 10**
PU type not supported by VTAM
- 11**
timer expired during VARY ACTIVATE
- 12**
error activating device
- 14**
device not defined to come online at IPL

15
device does not support sysdef automation

101
ISTDINFO internal error 101

102
ISTDINFO internal error 102

201
ISTDINFO internal error 201

301
ISTDINFO internal error 301

302
ISTDINFO internal error 302

303
ISTDINFO internal error 303

304
ISTDINFO internal error 304

305
ISTDINFO internal error 305

306
ISTDINFO internal error 306

307
ISTDINFO internal error 307

308
ISTDINFO internal error 308

309
ISTDINFO internal error 309

310
ISTDINFO internal error 310

311
ISTDINFO internal error 311

401
ISTDINFO internal error 401

501
ISTDINFO internal error 501

502
ISTDINFO internal error 502

503
ISTDINFO internal error 503

504
ISTDINFO internal error 504

505
ISTDINFO internal error 505

601
ISTDINFO internal error 601

602
ISTDINFO internal error 602

603
ISTDINFO internal error 603

604
ISTDINFO internal error 604

- 605**
ISTDINFO internal error 605
- 606**
ISTDINFO internal error 606
- 607**
ISTDINFO internal error 607
- 608**
ISTDINFO internal error 608
- 609**
ISTDINFO internal error 609
- 610**
ISTDINFO internal error 610
- 611**
ISTDINFO internal error 611

Appendix G. Message translation using the MVS Message Service

Multicultural support uses the MVS Message Service (MMS) for the translation of VTAM messages. This section describes how MMS is used to translate VTAM messages. For an overview of multicultural support, see [“Multicultural support” on page 199](#). This section contains information about the following topics:

- Overview of MMS support
- Internal translation
- External translation
- Skeleton file use

Overview of MMS support

The MVS Message Service translates VTAM messages using message skeleton files for particular languages. MMS finds a skeleton file matching the language requested and then searches for the message and returns the text from the skeleton file.

VTAM provides a base U.S. English message skeleton file named ISTVTMEU. This file contains all VTAM messages including USS, TSO/VTAM, and logon manager messages. The way in which ISTVTMEU is used depends on whether the translation of the message occurs internally or externally to VTAM.

Internal translation

VTAM uses internal translation to translate USS and TSO/VTAM user messages. These messages must be translated internal to VTAM so that they can be sent on either an SSCP-LU (for USS) or LU-LU session (for TSO/VTAM). For user messages, ISTVTMEU defines the U.S. English text that should be used for these messages. This text can be changed, providing a function similar to USS tables.

Selecting internal translation

To use internal translation:

1. To support a language other than U.S. English, first define a language skeleton file to MMS. The language skeleton file must have the language code or language name that will be specified by the logical unit. If the required language skeleton file does not exist when the USS command is entered and the MMS is active, the USS command will be rejected. For TSO/VTAM, default U.S. English messages are issued if the required language skeleton file does not exist.
2. An LU must explicitly request that translation be performed before messages will be translated internally.
 - For USS, the LANG operand on USS commands determines the language.
 - For TSO/VTAM, the language can be determined from the LANG or LANGTAB operands (the language information is provided on the CINIT) or by the TSO PROFILE command with the PLANG operand.

Defining messages for internal translation

The IBM-supplied file, ISTVTMEU, contains messages for internal and external translation. The user messages defined in ISTVTMEU for internal translation are USSMSG01 through USSMSG13 for USS and IKT00201I through IKT00405I for TSO/VTAM. For example, USS message 1 is defined as follows:

```
USSMSG01          INVALID &1. COMMAND SYNTAX    &TIME.
```

You can change the wording of this message to different wording if required. The &1. is an example of a substitution token and is used here to identify where the incorrect command name is to be placed in the message. Additional substitution tokens are &2., &3. and so on. Also, the &TIME., &DATE., &LUNAME., and &NETID. substitution tokens are used for the time, date, LU name, and network ID. If you want to provide USS messages in other languages using MMS, another skeleton file can be created for the required language with the required message text. For example, you could code the following in a new file named ISTVTFRA with the text in French:

```
USSMSG01      THE SYNTAX FOR COMMAND &1. IS INVALID
```

Then, if the user specified LANG(FRA) on a USS command, the above message text would be provided in the case of a USS command syntax error. The above message could also be defined as a multiline message, and, when translation occurs, all of the lines in the translated message are written to the LU. There is a 23-line limit on the number of lines in a translated message for USS and an 8-line limit for TSO/VTAM.

USS messages defined with MMS skeleton files must follow the same restrictions as USS messages defined with USS tables. Only a limited character set is available to SNA logical units on the SSCP-LU session.

For TSO/VTAM messages, any character set supported by the user terminal can be used, including double byte-character sets. Using incompatible character sets and code pages can result in erroneous output or I/O errors.

External translation

External translation occurs outside VTAM and is only available to authorized TSO/E 2.2 users using multiple console support (MCS) extended consoles.

Selecting external translation

In external translation, the language for the extended console is determined by the user profile, which is set with the PLANG operand of the TSO PROFILE command.

When a language other than U.S. English is selected, the following condition occurs:

1. ISTVTMEU is used to parse the U.S. English message to determine the variable text.
2. MMS looks up the corresponding translated message and substitutes the variable text.

ISTVTMEU and ISTINCNO must match for external translation to work correctly. It is recommended that you use the IBM-supplied copies of these files and that you do not use the USSTAB start option. Customization of messages can then be accomplished by changing the language skeleton files used by MMS.

Defining messages for external translation

Messages for external translation are also defined in ISTVTMEU. All of the messages for external translation are messages that are written to the system console. You should not change the U.S. English form of these messages in ISTVTMEU, because ISTVTMEU must correspond exactly with ISTINCNO for external translation to occur. For operator messages, you should use the IBM-supplied default table, ISTINCNO, rather than using a customized operator USS table.

For translation to other languages, provide a language skeleton file similar to ISTVTMEU with the translated text. For example, VTAM message IST001I is defined as follows in ISTVTMEU:

```
IST001I      001      VTAM START REJECTED - &1.
IST001I      002      &modnm. VTAM START REJECTED - &1.
```

The text of this message could be changed to French in a new skeleton file named ISTVTFRA. It is necessary to have two formats of this message: one for the message without the module name and one for the message with the module name. MMS will select the format that most closely matches

the message received and perform the translation by selecting the corresponding message with the corresponding format number from the new French skeleton file.

For example, assume ISTVTFRA has the following definition for IST001I:

```
IST001I      001      &1. - REJECT START
IST001I      002      &modnm. &1. - REJECT START
```

With the above skeletons coded, a TSO/E 2.2 user using an MCS extended console with the primary language of FRA will have the following VTAM message translated as shown below:

```
Message issued by VTAM:
IST001I VTAM START REJECTED - TERMINATION IN PROGRESS

Translated message on MCS Extended console
IST001I TERMINATION IN PROGRESS - REJECT START
```

The translated messages can be in any language supported by MMS, including double-byte character set languages. The only restriction is that the terminal being used for a given language must support that language with the correct character set and code page, or erroneous output or I/O errors can occur. Translated messages can also be defined as multiline messages if more than one line is needed to express the error in the translation language.

Skeleton file use

ISTVTMEU, and any other translation language skeletons required by the installation, must be compiled by the MMS compiler before use. Also, MMS must be active before VTAM can use any of the skeleton files.

ISTVTMEU is shipped in library SYS1.SISTDAT2. This member can be copied to another partitioned data set where other MMS skeleton files are kept. Skeleton files for languages other than U.S. English must be provided by the installation. For information about compiling message skeleton files, see [z/OS MVS Planning: Operations](#). For information about defining message skeleton files, see the [z/OS MVS Programming: Assembler Services Guide](#).

Appendix H. Forcing an APPN route in a VTAM network

This appendix describes how to force the path of a route through your network. You might want to force a route if you want certain sessions to use a specific path rather than the default path.

If you use only the defaults, the nodes and links are considered equal and the route with the fewest nodes and links will have the least weight. The route with the least weight is selected as the route for the session.

For example, if all the defaults are used in the network shown in [Figure 155 on page 613](#) and an LU in NN1 requests a session with an LU in NN3, the route chosen for the session will go from NN1 through NN2 to NN3.




Figure 155. Sample network showing default route

If you want to force a different route for a certain type of session, you need to cause that route to have a lower weight than the original default route. For example, you might want the session between the LU in NN1 and the LU in NN3 in [Figure 155 on page 613](#) to be routed through NN4 and NN5. By defining a new APPN Class of Service and a transmission group (TG) profile, you can create a lower weight for the route you choose than the weight of the default route.

This example demonstrates using the UPARM1 operand to force a route to traverse NN1, NN4, NN5, and NN3. With the defaults, the values for UPARM1 for the links in the network would look like the example in [Figure 156 on page 614](#).




Figure 156. Sample network using default for UPARM1

To create a lower weight for the chosen route, you can define a TG profile with a UPARM1 value of 0 and specify that TG profile on the PUs associated with the links between NN1 and NN4, NN4 and NN5, and NN5 and NN3. The UPARM1 values in the network would then appear as they do in [Figure 157](#) on page 614.




Figure 157. Sample network using TG profile on some links

By defining UPARM1 on the TG profile as 0, the links that specify that TG profile have a unique characteristic. You can define a new Class of Service that will give preference to links with a lower UPARM1 value. The following steps describe how to define the TG profile and define a Class of Service to take advantage of the UPARM1 value:

1. Code a TG profile with UPARM1=0 as shown in the following example.

```
*****
* TG Profile to change UPARM1 value                                     *
*****
*
UPARMLOW TGP    UPARM1=0
```

2. Specify this TG profile for the PUs associated with each line in the preferred route.
3. Code a Class of Service (COS) that gives preference to links using the TG profile.

Note: Use the APPNCOS statement to name the new COS UPARMCOS.

- a. Copy the eighth LINEROW and NODEROW statements from the IBM-supplied COS definition #CONNECT (found in the COSAPPN member of SYS1.VTAMLST).
- b. Make a second copy of the LINEROW statement.
- c. On the LINEROW statements, modify only the NUMBER, UPARM1, and WEIGHT operands.

Modify the first LINEROW statement to specify NUMBER=1, UPARM1=(0,99), and WEIGHT=30.

Modify the second LINEROW statement to specify NUMBER=2, UPARM1=(100,255), and WEIGHT=240.

4. File this new APPN COS definition in SYS1.VTAMLST under the new member, COSSAMP. Use the VARY ACT command to activate this definition by specifying ID=COSSAMP. (As an alternative, if VTAM is not already active, you can specify COSSAMP in your CONFIG list so that the definition will be activated as part of VTAM initialization.)
5. Define a mode table entry that specifies the new APPN COS definition and add it to ISTINCLM or the mode tables associated with the LUs (if the default mode table is not being used). The following code is an example of the mode table entry.

```
          TITLE 'UPARMODE'
*****
*
*          LOGMODE ENTRY FOR FORCED APPN ROUTE EXAMPLE -
*
*****
UPARMODE MODEENT LOGMODE=UPARMODE,FMPROF=X'03',
               TSPROF=X'03',PRIPROT=X'B1',SECPROT=X'A0',
               COMPROT=X'3040',
               APPNCOS=UPARMCOS
```

6. Specify the new mode name, UPARMODE, when requesting the session.
7. Because UPARMCOS is the APPN COS definition used by the route selection process, the NN1 to NN2 to NN3 route has a weight of $240 + 160 + 240 = 640$ based on the values of the WEIGHT operand from the UPARMCOS definition. Because the NN1 to NN4 to NN5 to NN3 route has a weight of $30 + 160 + 30 + 160 + 30 = 410$ based on the values of the WEIGHT operand, it will be the chosen least-weight route.

Note: For this example, it is required that the new mode and Class of Service tables be stored in the nodes that are the endpoints of the session path; however, we strongly recommend that the Class of Service and mode tables be consistent among all the nodes in the network.

The following example shows the new user-defined Class of Service.

```
*****
* user-defined class-of-service definition                             *
*****
UPARMCOS APPNCOS  PRIORITY=MEDIUM      transmission priority
          LINEROW  WEIGHT=30,           TG weight
          NUMBER=1, line row number
          UPARM1=(0,99), user defined parameter 1
```

UPARM2=(0,255),	user defined parameter 2	*
UPARM3=(0,255),	user defined parameter 3	*
CAPACITY=(MINIMUM,MAXIMUM),	line speed	*
COSTTIME=(0,255),	cost per connect time	*
COSTBYTE=(0,255),	cost per byte transmitted	*
PDELAY=(MINIMUM,MAXIMUM),	propagation delay	*
SECURITY=(UNSECURE,MAXIMUM)	security level for TG	
LINEROW WEIGHT=240,	TG weight	*
NUMBER=2,	line row number	*
UPARM1=(100,255),	user defined parameter 1	*
UPARM2=(0,255),	user defined parameter 2	*
UPARM3=(0,255),	user defined parameter 3	*
CAPACITY=(MINIMUM,MAXIMUM),	line speed	*
COSTTIME=(0,255),	cost per connect time	*
COSTBYTE=(0,255),	cost per byte transmitted	*
PDELAY=(MINIMUM,MAXIMUM),	propagation delay	*
SECURITY=(UNSECURE,MAXIMUM)	security level for TG	
NODEROW NUMBER=2,	node row number	*
WEIGHT=160,	node weight	*
CONGEST=(LOW,HIGH),	congestion	*
ROUTERES=(0,255)	route addition resistance	

Appendix I. Border node connection types

This appendix lists connection types for various VTAM and partner node combinations. For more information about border nodes and APPN multiple network connectivity support, see [“APPN multiple network connectivity”](#) on page 75.

Table 73. Connection type for selected VTAM and partner node combinations								
Local VTAM characteristics								
Partner node characteristics	BN		NATIVE			NETID		Connection type
	YES	NO or Not Coded	YES	NO	Not Coded	Same	Different	
Network node	X		X				X	Native connection, CP-CP sessions supported
Network node	X		X				X	Native connection, CP-CP sessions not supported
Network node	X			X		X		Peripheral connection, CP-CP sessions supported
Network node	X			X			X	Peripheral connection, CP-CP sessions supported
Network node	X				X	X		Native connection, CP-CP sessions supported
Network node	X				X		X	Peripheral connection, CP-CP sessions supported
Network node		X			X	X		Native connection, CP-CP sessions supported
Network node		X			X		X	Native connection, CP-CP sessions not supported
Peripheral border node	X		X			X		Native connection, CP-CP sessions supported
Peripheral border node	X		X				X	Connection failed by local VTAM

Table 73. Connection type for selected VTAM and partner node combinations (continued)

Local VTAM characteristics								
Partner node characteristics	BN		NATIVE			NETID		Connection type
	YES	NO or Not Coded	YES	NO	Not Coded	Same	Different	
Peripheral border node	X			X		X		Peripheral connection, CP-CP sessions supported
Peripheral border node	X			X			X	Peripheral connection, CP-CP sessions supported
Peripheral border node	X				X	X		Native connection, CP-CP sessions supported
Peripheral border node	X				X		X	Peripheral connection, CP-CP sessions supported
Peripheral border node		X			X	X		Native connection, CP-CP sessions supported
Peripheral border node		X			X		X	Peripheral connection (managed by partner node), CP-CP sessions supported
Extended border node, NATIVE=YES	X		X			X		Native connection, CP-CP sessions supported
Extended border node, NATIVE=YES	X		X				X	Native connection, CP-CP sessions not supported
Extended border node, NATIVE=YES	X			X		X		Connection failed by partner node
Extended border node, NATIVE=YES	X			X			X	Connection failed by partner node
Extended border node, NATIVE=YES	X				X	X		Native connection, CP-CP sessions supported
Extended border node, NATIVE=YES	X				X		X	Connection failed by partner node

Table 73. Connection type for selected VTAM and partner node combinations (continued)

Local VTAM characteristics								
Partner node characteristics	BN		NATIVE			NETID		Connection type
	YES	NO or Not Coded	YES	NO	Not Coded	Same	Different	
Extended border node, NATIVE=YES		X			X	X		Native connection, CP-CP sessions supported
Extended border node, NATIVE=YES		X			X		X	Native connection, CP-CP sessions not supported
Extended border node, NATIVE=NO	X		X			X		Connection failed by local VTAM
Extended border node, NATIVE=NO	X		X				X	Connection failed by local VTAM
Extended border node, NATIVE=NO	X			X		X		Extended connection, CP-CP sessions supported
Extended border node, NATIVE=NO	X			X			X	Extended connection, CP-CP sessions supported
Extended border node, NATIVE=NO	X				X	X		Extended connection, CP-CP sessions supported
Extended border node, NATIVE=NO	X				X		X	Extended connection, CP-CP sessions supported
Extended border node, NATIVE=NO		X			X	X		Peripheral connection (managed by partner node), CP-CP sessions supported
Extended border node, NATIVE=NO		X			X		X	Peripheral connection (managed by partner node), CP-CP sessions supported
Extended border node, NATIVE not coded	X		X			X		Native connection, CP-CP sessions supported

Table 73. Connection type for selected VTAM and partner node combinations (continued)

Local VTAM characteristics								
Partner node characteristics	BN		NATIVE			NETID		Connection type
	YES	NO or Not Coded	YES	NO	Not Coded	Same	Different	
Extended border node, NATIVE not coded	X		X				X	Connection failed by partner node
Extended border node, NATIVE not coded	X			X		X		Extended connection, CP-CP sessions supported
Extended border node, NATIVE not coded	X			X			X	Extended connection, CP-CP sessions supported
Extended border node, NATIVE not coded	X				X	X		Native connection, CP-CP sessions supported
Extended border node, NATIVE not coded	X				X		X	Extended connection, CP-CP sessions supported
Extended border node, NATIVE not coded		X			X	X		Native connection, CP-CP sessions supported
Extended border node, NATIVE not coded		X			X		X	Peripheral connection (managed by partner node), CP-CP sessions supported

Notes:

1. The local VTAM node is always configured with the NODETYPE=NN specified.
2. XNETALS=YES is assumed whenever the NETIDs of the two nodes are different. If XNETALS=NO is specified, connections between two nodes with different NETIDs fail.
3. When BN=YES is coded, XNETALS=YES is the default start option. When BN=YES is not coded, XNETALS=NO is the default start option.

Appendix J. VTAM restricted materials

The VTAM modules and macroinstructions listed in this appendix are considered to be Restricted Materials of IBM and are available through the View Program Listings (VPL) application. This information can be accessed by using ServiceLink or Dial IBM. All other VTAM modules are designated as object code only (OCO).

Table 74. User replaceable or modifiable modules

Module	Description
ISTMGC01	CNM routing table (IBM-provided default table)
ISTINCLM	Logon mode table
ISTINCDT	Session-level USS table

Table 75. VTAM message modules

Module	Description
ISTCFCMM	English text of VTAM messages
ISTINCNO	Formatted English text of VTAM messages
ISTCFCML	Defines variables and variable lengths for VTAM messages

Appendix K. 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)

The following RFC also contains SNA architectural specifications:

- RFC 2353 *APPN/HPR in IP Networks APPN Implementers' Workshop Closed Pages Document*

RFCs are available at <http://www.rfc-editor.org/rfc.html>.

Appendix L. 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)

The following RFC also contains SNA architectural specifications:

- RFC 2353 *APPN/HPR in IP Networks APPN Implementers' Workshop Closed Pages Document*

RFCs are available at <http://www.rfc-editor.org/rfc.html>.

Appendix M. Accessibility

Accessible publications for this product are offered through [IBM Documentation for z/OS](#).

If you experience difficulty with the accessibility of any z/OS documentation see [How to Send Feedback to IBM](#) to leave documentation feedback.

Notices

This information was developed for products and services that are offered in the USA or elsewhere.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 United States of America

For license inquiries regarding double-byte character set (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.

This information could include missing, incorrect, or broken hyperlinks. Hyperlinks are maintained in only the HTML plug-in output for the IBM Documentation. Use of hyperlinks in other output formats of this information is at your own risk.

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 wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation Site Counsel 2455 South Road Poughkeepsie, NY 12601-5400 USA

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 document 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 measurements 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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate 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.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's name, email address, phone number, or other personally identifiable information for purposes of enhanced user usability and single sign-on configuration. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at ibm.com/privacy and IBM's Online Privacy Statement at ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at ibm.com/software/info/product-privacy.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, 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.

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: [IBM Lifecycle Support for z/OS \(www.ibm.com/software/support/systemsz/lifecycle\)](http://www.ibm.com/software/support/systemsz/lifecycle)
- For information about currently-supported IBM hardware, contact your IBM representative.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, HCD, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at [Copyright and trademark information at www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Bibliography

This bibliography contains descriptions of the documents in the z/OS Communications Server library.

z/OS Communications Server documentation is available online at the z/OS Internet Library web page at <http://www.ibm.com/systems/z/os/zos/library/bkserv/>.

z/OS Communications Server library updates

Updates to documents are also available on RETAIN and in information APARs (info APARs). Go to <https://www.ibm.com/mysupport> to view information APARs.

- [z/OS Communications Server V2R1 New Function APAR Summary](#)
- [z/OS Communications Server V2R2 New Function APAR Summary](#)
- [z/OS Communications Server V2R3 New Function APAR Summary](#)
- [z/OS Communications Server V2R4 New Function APAR Summary](#)

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 functions, 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 Appl 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 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	<p>This document enables you to customize SNA, and includes the following information:</p> <ul style="list-style-type: none"> • Communication network management (CNM) routing table • Logon-interpret routine requirements • Logon manager installation-wide exit routine for the CLU search exit • TSO/SNA installation-wide exit routines • SNA installation-wide exit routines

Writing application programs

Title	Number	Description
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.

Title	Number	Description
z/OS Communications Server: CSM Guide	SC27-3647	This document describes how applications use the communications storage manager.

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

2 [184](#), [472](#)
2.1, nonnative network connection [184](#), [472](#)
31-bit backed data space [581](#)
3172 Nways Interconnect Controller
 APPN connections [48](#)
 connection network, defining a connection to [53](#)
 tuning statistics [503](#)
 VCNS and [329](#)
3270 IDS
 configuration [307](#)
 considerations and assessment [301](#)
 deployment strategy [307](#)
 environment [301](#), [302](#)
 exploitation cost
 system resource cost [306](#)
 incident [311](#)
 known application 3270 solutions [307](#)
 overview [299](#)
 SNA technologies
 environmental factors [304](#)
 network connectivity [304](#)
 VTAM commands [309](#)
3270 IDS incident [311](#), [315](#)
3270 Intrusion Detection Services
 considerations and assessment [301](#)
 overview [299](#)
3272 cluster controllers [185](#)
6.2 sessions
 description [335](#)
 non-LU 6.2 sessions and [336](#)
64-bit backed data space [581](#)

A

ACB (access method control block), opening by application program [277](#)
ACB macroinstruction
 APPLID operand [277](#)
 PERSIST operand [327](#)
ACBNAME operand on APPL definition statement [277](#)
accessibility
 contact IBM [627](#)
accounting for cross-domain sessions [217](#)
activating resources
 ADJSSCP tables [482](#)
 application programs [486](#)
 cross-domain resource managers [482](#)
 multiple-domain environment [482](#)
 multiple-network environment [483](#)
 order of activation [482](#)
adding PUs and LUs dynamically [174](#)
address space identifier (ASID) [583](#)
address space, spanning multiple [329](#)
address subarea node
 address incompatibility [268](#)

address subarea node (*continued*)
 addressing methods [267](#)
 extended addressing compatibility [270](#)
 extended network addressing [267](#)
 path definition, subarea [266](#)
 subarea addressing restrictions [270](#)
address translation [463](#)
address, element [267](#)
addressing
 compatibility
 extended network and extended subarea [270](#)
 nonextended and extended network [267](#)
 LU, dependent [196](#)
 multiple-network (SNI) [459](#)
 overview [267](#)
 restrictions [268](#)
 structure types [267](#)
 subarea restrictions [270](#)
 translation of address [463](#)
addressing, enhanced for session managers [26](#)
adjacent link station (ALS) [189](#)
adjacent SSCP
 lists for CDRSCs [431](#)
 multiple-domain environment [431](#)
 overview [431](#)
 performance [435](#)
 request routing [457](#)
 selecting with the session management exit [216](#)
 tables
 alias name translation and [458](#)
 combined APPN and subarea network [413](#)
 default lists [454](#)
 defining [432](#), [454](#)
 dynamic CDRSCs and [458](#)
 dynamically defining [432](#)
 multiple-network environment [454](#)
 multiple-network sample [454](#)
 network-specific lists [454](#)
 requirements for [454](#)
 types of [454](#)
adjacent SSCP lists for CDRSCs [431](#)
adjacent SSCP table, dynamic
 CDRSC definition statement [427](#)
 controlled by another VTAM domain [427](#)
 cross-domain resource major node [427](#)
 dynamic adjacent SSCP table [432](#)
 locating cross-domain resources [431](#)
 network resources for XRF [428](#)
 USERVAR with XRF [428](#)
ADJCP definition statement [388](#)
alerts, VTAM to NPDA, NMVTLOG start option [27](#)
alias name translation
 adjacent SSCP tables and [458](#)
 alias selection function (VTAM) [465](#)
 defining [467](#)
 LU 6.2 [470](#)
 NetView [474](#)

alias name translation (*continued*)

overview [466](#)

allocating private storage [329](#)

ALS (adjacent link station) [189](#)

ALSLIST operand

APPN connections [386](#)

combined networks [413](#)

independent LUs [193](#)

ALSREQ start option [193](#), [195](#)

ANR (automatic network routing) [390](#)

APBUF buffer pool

number of buffers per page [542](#)

overview [536](#)

APF (authorized program facility) [284](#)

API, cross-memory [329](#)

APING support [38](#), [337](#), [488](#)

APPC operand on APPL definition statement [330](#)

APPCMD macroinstruction

CONFTXT operand [299](#)

requirement for LU 6.2 [330](#)

APPL definition statement

ACBNAME operand [277](#)

APPC operand [330](#)

ATNLOSS operand [335](#)

AUTH operand [283](#)

AUTOSES operand [333](#)

DDRAINL operand [333](#)

description [277](#)

DLOGMODE operand [318](#)

DMINWNL operand [333](#)

DMINWNR operand [333](#)

DRESPL operand [333](#)

DSESLIM operand [333](#)

EAS operand [528](#)

ENCR operand [290](#)

HAVAIL operand [325](#)

LIMQSINT operand [336](#)

LIMRES operand [336](#)

MODETAB operand [318](#)

PARSESS operand [284](#), [464](#)

sample [277](#)

SYNCLVL operand [335](#)

VCNS operand [329](#)

APPL EAS storage estimates [579](#)

application password [590](#)

application program

ACBNAME operand [277](#)

acquiring sessions [284](#)

activating [486](#)

APPLID operand [277](#)

ATTN exit [335](#)

authorized path [284](#)

communication network management (CNM) [495](#)

confidential data [299](#)

controlling [210](#)

cross-memory API [329](#)

data transfer, privileged [284](#)

defining [277](#)

dial number digits, overriding [283](#)

draining conversations [333](#)

encryption [289](#)

extended recovery facility (XRF) [325](#)

failed sessions, reconnecting to [327](#)

generic resources [324](#)

application program (*continued*)

impact of outages, reducing [325](#)

logon mode [318](#)

logon mode parameters [318](#)

logon mode table entries, create [318](#)

logon request validation [283](#)

LU 6.2 [330](#)

LU 6.2 security [334](#)

LU 6.2 session [330](#)

LU 6.2 sync point service [335](#)

model definition [278](#)

naming [277](#)

network management capability [494](#)

overview [19](#)

parallel sessions [284](#), [464](#)

passing logon request [283](#)

private storage, allocating [329](#)

program operator (PPO and SPO) [494](#)

resource verification reduction [331](#)

sample definition [277](#)

security facilities [288](#)

security protocol, conversation [334](#)

spanning multiple address space [329](#)

start-stop devices, communicating with [329](#)

starting

name of application and [277](#)

overview [298](#)

TSO [557](#)

USERVAR [319](#)

validate logon [283](#)

XRF (extended recovery facility) [325](#)

application-supplied operands for dial [283](#)

APPLID operand on ACB macroinstruction [277](#)

APPN

addressing overview [16](#)

ALSLIST operand on CDRSC definition statement [386](#)

bisynchronous (BSC) 3270 sessions [405](#)

border node

connection types [617](#)

description [6](#), [75](#)

BSC 3270 sessions [405](#)

central directory server

benefits of [238](#)

CDSREFER start option [239](#)

overview [6](#)

resource registration [239](#)

search requests [239](#)

using to improve efficiency [238](#)

Class of Service (COS), overview [15](#), [242](#)

combined with subarea network [401](#)

composite network node, overview [8](#)

connection network [50](#)

connections

3172 Nways interconnect controller [48](#)

assigning NCP addresses [386](#)

converting LEN connections to APPN [401](#)

definitions of [385](#)

leased lines [47](#)

overview [39](#)

requirements [385](#)

sample type 2.1 channel [47](#)

session control block usage [386](#)

setting maximum number of sessions [386](#)

type 2.1 channel [40](#), [402](#)

APPN (continued)

- control point (CP)
 - CP name [22](#), [387](#)
 - overview [4](#)
- converting LEN connections [401](#), [402](#)
- CP name for VTAM [22](#)
- CPNAME operand on PU definition statement [386](#)
- dependent LUs [405](#)
- directed search [236](#)
- directory services, controlling the size of [240](#)
- end node, overview [5](#)
- extended subnetwork boundaries [77](#)
- generic resources [324](#)
- host-to-host channel connection [42](#)
- IDBLK operand on PU definition statement [386](#)
- IDNUM operand on PU definition statement [386](#)
- implementing a network [385](#)
- interchange node
 - coding [402](#)
 - combined APPN and subarea network [402](#)
 - overview [7](#)
- LEN connections, converting [401](#), [402](#)
- locating nodes [235](#), [404](#)
- migrating from subarea [401](#)
- migration data host
 - in APPN [403](#)
 - overview [8](#)
- multiple network connectivity [75](#)
- network accessible units [9](#)
- network node server
 - defining list [85](#)
 - overview [5](#)
 - replacing list [86](#)
 - resource registration [239](#)
 - sample list [85](#)
- network node, overview [4](#)
- nodes
 - adding and moving [420](#)
 - adjacent [387](#)
 - connecting to VTAM [39](#)
 - types of [3](#)
- peripheral subnetwork boundary [77](#)
- resource registration [239](#)
- resource verification reduction [331](#)
- route calculation, overview [16](#)
- route selection, overview [16](#)
- routing
 - directory services [235](#)
 - improving efficiency [238](#)
 - route selection [16](#)
 - topology database [14](#)
 - types of searches [236](#)
- searches
 - broadcast search [236](#)
 - directed search [236](#)
 - improving efficiency [238](#)
 - resource verification reduction [331](#)
- sessions in a combined network [412](#), [472](#)
- start options
 - APPNCOS [79](#)
 - BN [77](#)
 - BNDYN [78](#)
 - BNORD [78](#)
 - CDSERVER [239](#)

APPN (continued)

- start options (continued)
 - CDSREFER [239](#)
 - CONNTYPE [401](#)
 - CPCP [387](#)
 - DIRSIZE [240](#)
 - DIRTIME [240](#)
 - DYNADJCP [387](#)
 - HPR [391](#)
 - HRPST [391](#)
 - INITDB [242](#)
 - NODETYPE [402](#), [403](#)
 - SNVC [78](#)
 - VERIFYCP [334](#)
 - VFYRED [332](#)
 - VFYREDTI [332](#)
 - VRTG [82](#)
 - VRTGCPCP [82](#)
 - XID operand on PU definition statement [385](#)
- APPNCOS start option [79](#)
- APPNTOSA COS mapping table [420](#)
- ASID (address space identifier) [583](#)
- ASIRFMSG start option [26](#)
- assistive technologies [627](#)
- associated LU table [202](#)
- Asynchronous transfer mode connections
 - ATM LAN emulation connections [55](#)
 - ATM native connections [55](#)
- ATCCON (configuration list) [28](#)
- ATCSTR00 (start option list) [20](#), [22](#), [183](#)
- ATM connections
 - ATM LAN emulation connections [55](#)
 - ATM native connections [55](#)
- ATM native connections
 - connection to the Open Systems Adapter [56](#)
 - port on the Open Systems Adapter [58](#)
 - TGs over SVCs [62](#)
- ATNLOSS operand on APPL definition statement [335](#)
- ATTN exit [335](#)
- AUTH operand on APPL definition statement [283](#)
- AUTHLEN start option and operand [196](#)
- authorized path [284](#)
- authorized program facility (APF) [284](#)
- authorized transmission priority for LEN connections [196](#)
- authorizing cross-domain sessions [216](#)
- automatic logon
 - changing an LU's automatic logon specification [212](#)
 - coding [209](#)
 - controlling [216](#)
 - defined [209](#)
 - fails, warning [323](#)
 - LOGAPPL operand [209](#), [210](#)
 - multiple-network [473](#)
 - operator-controlled [212](#)
- automatic network routing (ANR) [390](#)
- AUTORTRY start option [212](#)
- AUTOSER operand on APPL definition statement [333](#)
- AUTOTI start option [212](#)

B

- backup
 - sessions, extended recovery facility and [325](#)
- backup start option list (LISTBKUP) [23](#)

- binary synchronous (BSC) connection and APPN [405](#)
- BIND, extended [189](#)
- BN start option [77](#)
- BNDYN start option [78](#)
- BNORD start option [78](#)
- border node
 - connection types [617](#)
 - description [6](#), [75](#)
- boundary function-based transmission groups [39](#)
- broadcast search (APPN)
 - central directory server and [237](#)
 - limiting [238](#), [414](#)
 - originating [236](#)
 - overview [236](#)
 - when used [237](#)
- BSBUF buffer pool
 - number of buffers per page [542](#)
 - overview [536](#)
- BSC (binary synchronous) connection and APPN [405](#)
- BSCMDRS start option [26](#)
- buffer pools
 - allocating storage [539](#)
 - dynamic expansion
 - basic allocation and [539](#)
 - guidelines for [541](#)
 - operation of [539](#)
 - virtual storage and [543](#)
 - monitoring [500](#)
 - number of buffers per page [543](#)
 - overview [500](#)
 - types of [536](#)
- buffer pools, CSM
 - size of [581](#)
 - source of [581](#)

C

- CACHETI start option [79](#)
- canceling VTAM [494](#)
- CDRM (cross-domain resource manager)
 - activating [482](#)
 - coding of [424](#)
 - deactivating [492](#)
 - description of [423](#)
 - DISJOINT operand [414](#)
 - dynamic definition of cross-domain resources [425](#)
 - external CDRM [423](#)
 - gateway VTAM, for [444](#)
 - host CDRM [423](#)
 - multiple domains [423](#)
 - sample of [424](#)
- CDRSC (cross-domain resource)
 - adjacent SSCP (ADJSSCP) table, defining [454](#)
 - adjacent SSCP lists for CDRSCs [431](#)
 - CDRM ownership
 - changing by operator [438](#)
 - changing dynamically [437](#)
 - verifying [437](#)
 - combined APPN and subarea network [410](#)
 - defining [410](#), [425](#)
 - definition statement, VFYOWNER operand [437](#)
 - description [454](#)

- CDRSC (cross-domain resource) (*continued*)
 - dynamic [452](#)
 - locating
 - adjacent SSCP table [431](#)
 - ADJCDRM definition statement [432](#)
 - CDRSC routed to each SSCP [431](#)
 - cross-network [454](#)
 - multiple-network configuration [455](#)
 - performance improving [435](#)
 - predefined [450](#)
 - SSCPDYN start option [435](#)
 - SSCPORD start option [435](#)
 - CDRSCTI start option [404](#), [414](#)
 - CDSERVER start option (APPN) [239](#)
 - CDSREFER start option (APPN) [239](#)
 - central directory server (APPN)
 - benefits of [238](#)
 - broadcast search and [237](#)
 - defining [239](#)
 - overview [6](#)
 - resource registration [239](#)
 - search requests [239](#)
 - selection of [239](#)
 - using to improve efficiency [238](#)
 - change-number-of-sessions
 - contention loser and [333](#)
 - modifying parameters [332](#)
 - operating commands for [333](#)
 - overview [332](#)
 - channel
 - connections, VTAM-to-peripheral node [185](#)
 - connections, VTAM-to-VTAM [87](#)
 - device name
 - non-SNA devices, for [185](#)
 - SNA devices, for [186](#)
 - enterprise systems connection (ESCON) [187](#)
 - multipath channel (MPC) for hosts [40](#)
 - sample APPN [47](#)
 - type 2.1 APPN [40](#)
 - channel-attached
 - APPN type 2.1 [40](#)
 - between host [87](#)
 - I/O performance with nodes [503](#)
 - loop-adapter attachment [187](#)
 - MVS/CP devices [170](#)
 - non-SNA devices [185](#)
 - SNA cluster controllers, defining statically [185](#)
 - type 2.1 [184](#)
 - channel-to-channel adapter connections
 - defining [87](#)
 - tuning I/O for [512](#)
 - checkpointing TRS and directory databases [242](#)
 - CICS, extended recovery facility (XRF) [325](#)
 - Class of Service (COS)
 - APPN, overview [15](#)
 - combined APPN and subarea network [416](#)
 - database (APPN) [15](#)
 - defaults, IBM-specified [261](#)
 - logon mode table [318](#)
 - mapping table [420](#)
 - multiple-network [461](#)
 - name [467](#)
 - substituting parameters [261](#)

- Class of Service (COS) (*continued*)
 - table [260](#)
 - unnamed entry [260](#)
- cluster controllers
 - non-SNA, defining [185](#)
 - SNA
 - defining statically [185](#)
 - I/O performance to [503](#)
 - sample tuning statistics report [511](#)
- CMC (communication management configuration)
 - overview [7](#)
- CMPMIPS start option [285](#)
- CNM (communication network management)
 - application program [495](#)
 - application program, defined [495](#)
 - filtering session awareness data [495](#)
 - NetView program and [475](#)
 - procedure-related identifier (PRID) [495](#)
 - routing table [495](#)
 - routing unsolicited requests [495](#)
- CNOS
 - contention loser and [333](#)
 - modifying parameters [332](#)
 - operating commands for [333](#)
 - overview [332](#)
- coattailing
 - channel-to-channel adapter
 - DELAY operand [554](#)
 - I/O buffer size [553](#)
 - virtual route (VR) window size [554](#)
 - DELAY operand [547](#)
 - determining amount of [527](#)
 - inbound [547](#)
 - MAXBFRU operand [548](#)
 - maximizing [546](#)
 - outbound [547](#)
 - overview [500](#)
 - UNITSZ value [548](#)
- coded data
 - CDRM key [594](#)
 - cross-domain cryptography [594](#)
 - filing secondary logic unit (SLU) keys [593](#)
 - single-domain encryption [593](#)
- coding concepts [29](#)
- coding resource definition statements
 - overview [29](#)
 - sift-down [30](#)
- combined APPN and subarea network [401](#)
- common service area (CSA)
 - buffer pool waste [502](#)
 - common service area 24-bit (CSA24) [529](#)
 - CSA limit [528](#)
 - monitoring [500](#)
- communication controller
 - I/O performance with [503](#)
 - sample tuning statistics report [511](#)
 - tuning I/O for [504](#), [511](#)
- communication management configuration (CMC)
 - overview [7](#)
- communication network management (CNM)
 - application program [495](#)
 - application program, defined [495](#)
 - filtering session awareness data [495](#)
 - NetView program and [475](#)
- communication network management (CNM) (*continued*)
 - procedure-related identifier (PRID) [495](#)
 - routing table [495](#)
 - routing unsolicited requests [495](#)
- Communications Server for z/OS, online information [xxxiv](#)
- communications storage manager
 - definition of [581](#)
 - starting [582](#)
 - system definition [581](#)
- Communications storage manager [581](#)
- composite network node
 - overview [8](#)
 - routing failure message [253](#)
- compression of data [284](#)
- confidential data [299](#)
- CONFIG operand on START command [28](#)
- configuration list
 - advantages of [20](#)
 - CONFIG operand [28](#)
 - for gateway VTAM [442](#)
 - overview [19](#)
 - sample [28](#)
- configuration of I/O devices,
 - dynamic
 - command list programs [601](#)
 - customization [173](#), [325](#)
 - default naming convention [173](#)
 - installation and preparation [171](#)
 - overview [170](#)
- configuration restart
 - defined [477](#)
 - NODELST data set [478](#)
 - WARM start [478](#)
- configuration restart file, VSAM
 - associating with major nodes [479](#)
 - CONFGDS operand [479](#)
 - CONFGPW operand [479](#)
 - MVS [479](#)
- Configure
 - 3270 IDS [307](#)
- CONFTXT operand on APPCCMD macroinstruction [299](#)
- congestion, route [273](#)
- connecting networks [443](#)
- connection networks (APPN)
 - overview [50](#)
- CONNTYPE
 - operand on PU definition statement, overview [385](#)
 - parallel transmission groups and [40](#)
 - start option [385](#), [401](#)
- consoles, multiple operator [489](#)
- contact
 - z/OS [627](#)
- contention loser, CNOS and [333](#)
- control point (CP)
 - adjacent
 - ADJCP definition statement [388](#)
 - DYNADJCP start option [387](#)
 - predefinition of [388](#)
 - dynamic creation of adjacent CPs [387](#)
 - overview [4](#)
- control vector [29](#) [366](#)
- conversations, LU 6.2
 - draining [333](#)
 - overview [330](#)

- converting LEN connections to APPN [401](#)
- COS (Class of Service)
 - APPN, overview [15](#)
 - combined APPN and subarea network [416](#)
 - database (APPN) [15](#)
 - defaults, IBM-specified [261](#)
 - logon mode table [318](#)
 - mapping table [420](#)
 - multiple-network [461](#)
 - name [467](#)
 - substituting parameters [261](#)
 - table [260](#)
 - unnamed entry [260](#)
- coupling facility failures
 - Sysplex Wide Security Associations
 - TCP/IP stack [381](#)
 - VTAM node [381](#)
 - Sysplexports [383](#)
- Coupling facility failures
 - Sysplexports
 - TCP/IP stack [383](#)
 - VTAM node [383](#)
- coupling facility structure
 - allocating storage for [348](#)
 - attributes [344](#)
 - connecting to [348](#)
 - duplexing [349](#)
 - dynamic altering of [349](#)
 - failures for generic resource configuration [359](#)
 - failures for multinode persistent session configuration [373](#)
 - for multinode persistent sessions [369](#)
 - rebuild [348](#)
 - setting up the sysplex environment [341](#)
 - size, determining [346](#)
 - storage shortages [349](#)
- CP (control point)
 - adjacent
 - ADJCP definition statement [388](#)
 - DYNADJCP start option [387](#)
 - predefinition of [388](#)
 - dynamic creation of adjacent CPs [387](#)
 - overview [4](#)
- CP name for VTAM [22](#)
- CP-CP sessions (APPN)
 - establishing between VTAM nodes [387](#)
 - overview [11](#)
- CPCP, start option [387](#)
- CPNAME and PU name uniqueness [401](#)
- CPNAME operand on PU definition statement
 - type 2.1 PU and [183](#)
- CPSVRMGR pipe [406](#)
- CPSVRMGR session [405](#)
- cross-domain resource (CDRSC)
 - adjacent SSCP (ADJSSCP) table, defining [454](#)
 - adjacent SSCP lists for CDRSCs [431](#)
 - CDRM ownership
 - changing by operator [438](#)
 - changing dynamically [437](#)
 - verifying [437](#)
 - combined APPN and subarea network [410](#)
 - defining [410](#), [425](#)
 - definition statement, VFYOWNER operand [437](#)
 - description [454](#)
- cross-domain resource (CDRSC) (*continued*)
 - dynamic [452](#)
 - locating
 - adjacent SSCP table [431](#)
 - ADJCDRM definition statement [432](#)
 - CDRSC routed to each SSCP [431](#)
 - cross-network [454](#)
 - multiple-network configuration [455](#)
 - performance improving [435](#)
 - predefined [450](#)
 - SSCPDYN start option [435](#)
 - SSCPORD start option [435](#)
- cross-domain resource control
 - CDRSC definition statement [427](#)
 - controlled by another VTAM domain [427](#)
 - cross-domain resource major node [427](#)
 - dynamic adjacent SSCP table [432](#)
 - locating cross-domain resources [431](#)
 - network resources for XRF [428](#)
 - USERVAR with XRF [428](#)
- cross-domain resource manager (CDRM)
 - activating [482](#)
 - coding of [424](#)
 - deactivating [492](#)
 - description of [423](#)
 - DISJOINT operand [414](#)
 - dynamic definition of cross-domain resources [425](#)
 - external CDRM [423](#)
 - gateway VTAM, for [444](#)
 - host CDRM [423](#)
 - multiple domains [423](#)
 - sample of [424](#)
- cross-memory API [329](#)
- cross-network routing
 - communication between network [459](#)
 - COS tables [463](#)
 - extended subarea address gateway [459](#)
 - extended subarea addressing [459](#)
 - gateway NCP [459](#)
 - gateway VTAM [459](#)
 - virtual route, default [463](#)
- CRPLBUF buffer pool
 - guidelines for [541](#)
 - overview [536](#)
- cryptographic capability, changing [599](#)
- cryptographic key
 - CDRM key [594](#)
 - cross-domain cryptography [594](#)
 - filing secondary logic unit (SLU) keys [593](#)
 - single-domain encryption [593](#)
- cryptography
 - application program [289](#)
 - APPN and cryptography [294](#)
 - end-to-end cryptography [293](#)
 - extended recovery facility (XRF) [326](#)
 - host-by-host cryptography [295](#)
 - sample
 - multiple-domain [291](#)
 - single-domain [290](#)
- CSA (common service area)
 - buffer pool waste [502](#)
 - common service area 24-bit (CSA24) [529](#)
 - CSA limit [528](#)
 - monitoring [500](#)

- CSA24 start option [529](#)
- CSALIMIT start option [528](#)
- CSM
 - buffer pools [581](#)
 - definition of [581](#)
 - starting [582](#)
 - storage types [581](#)
 - system definition [581](#)
- CSM installation and definition [581](#)
- CSM messages [582](#)
- CSM parmlib member [581](#)
- CSM tracing [583](#)
- CTCA
 - defining [87](#)
 - tuning I/O for [512](#)
- CWALL threshold condition (NCP) [502](#)

D

- data compression [284](#)
- data control, subarea to network [271](#)
- data hosts, overview [7](#)
- data transmission shutdown [273](#)
- DDRAINL operand on APPL definition statement [333](#)
- deactivating resources
 - automatic deactivation of cross-subarea links [492](#)
 - cross-domain resource managers [492](#)
 - forced deactivation [492](#)
 - forced reactivation [493](#)
 - immediate deactivation [492](#)
 - multiple-domain environment, in a [491](#)
 - order of deactivation [491](#)
 - overview [491](#)
- DELAY operand on LINE definition statement
 - CTC coattailing and [554](#)
 - inbound coattailing and [547](#)
 - NCP coattailing and [549](#)
- deleting PUs and LUs dynamically [174](#)
- dependent LU requester [405](#)
- dependent LU server [405](#)
- dependent LUs in an APPN network [405](#)
- diagnosis
 - alerts to NPDA [27](#)
 - module name, viewing [27](#)
 - NMVTLOG start option [27](#)
 - SDLC statistical MDRs [28](#)
 - start options [26](#)
 - time interval notification [27](#)
 - trace facility [28](#)
- dial number digits, overriding [283](#)
- directed search (APPN)
 - originating [236](#)
 - overview [236](#)
 - resource verification reduction [331](#)
 - when used [236](#)
- directory services (APPN)
 - checkpointing [242](#)
 - controlling the size of the database [240](#)
 - database [235](#)
 - improving performance [242](#)
 - learning about resources [235](#)
 - overview [235](#)
- DIRSIZE start option (APPN) [240](#)

- DIRTIME start option (APPN) [240](#)
- DISJOINT operand on CDRM definition statement [414](#)
- DISPLAY APING command [38](#), [337](#), [488](#)
- DISPLAY BFRUSE operator command [500](#)
- DISPLAY CSM command [582](#)
- DISPLAY STORUSE pools [529](#)
- DLCADDR operand on PATH definition statement
 - dependent LU requester, defining [406](#)
- DLOGMODE operand on APPL definition statement [318](#)
- DLUR [405](#)
- DLURNAME operand on PATH definition statement [406](#)
- DLUS [405](#)
- DMINWNL operand on APPL definition statement [333](#)
- DMINWNR operand on APPL definition statement [333](#)
- DNS, online information [xxxv](#)
- DRESPL operand on APPL definition statement [333](#)
- DSESLIM operand on APPL definition statement [333](#)
- DSIRFMSG start option [26](#)
- DYNADJCP start option [387](#)
- dynamic definition of CDRSCs [425](#)
- dynamic definition of independent LUs [191](#)
- dynamic definition of PUs [167](#)
- dynamic definition of VTAM-to-VTAM connections [350](#)
- dynamic I/O configuration
 - command list programs [601](#)
 - customization [173](#), [325](#)
 - default naming convention [173](#)
 - installation and preparation [171](#)
 - overview [170](#)
- dynamic path update [274](#)
- dynamic reconfiguration
 - effect on WARM start [479](#)
 - overview [174](#)
 - VARY ACT, UPDATE technique [176](#)
 - VARY DRDS technique [178](#)
- dynamic resource definition (model name) [201](#)
- dynamic selection of session connections [195](#)
- dynamic switched definitions [167](#)
- dynamic switched major node [167](#)
- dynamic table replacement [491](#)
- DYNDLGMD start option [192](#), [204](#)
- DYNLU start option [387](#)
- DYNMODTB start option [192](#), [204](#), [417](#)
- DYNPU operand [167](#)

E

- EAS (estimated number of active sessions)
 - APPL definitions [528](#)
 - LFBUF and [542](#)
 - LU definitions [528](#)
 - TSO, specifying for [558](#)
- EAS storage estimates [579](#)
- element address [267](#)
- ENA (extended network addressing)
 - addressing compatibility [267](#)
 - addressing restrictions [268](#)
 - compatibility, nonextended address [267](#)
 - element address [267](#)
- ENCR operand
 - APPL definition statement [290](#)
 - LU definition statement [290](#)
- encryption facility
 - cryptographic keys, defining [593](#)

- encryption facility (*continued*)
 - description [289](#)
 - extended recovery facility (XRF) [326](#)
 - sample
 - multiple-domain [291](#)
 - single-domain [290](#)
- encryption of data
 - application program [289](#)
 - APPN and cryptography [294](#)
 - CDRM key [594](#)
 - cross-domain cryptography [594](#)
 - end-to-end cryptography [293](#)
 - extended recovery facility (XRF) [326](#)
 - filing secondary logic unit (SLU) keys [593](#)
 - host-by-host cryptography [295](#)
 - sample
 - multiple-domain [291](#)
 - single-domain [290](#)
 - single-domain encryption [593](#)
- encryption of data, changing [599](#)
- end nodes (APPN)
 - coding [403](#)
 - combined APPN and subarea network [403](#)
 - overview [5](#)
 - resource registration [239](#)
- ENHADDR start option [25](#)
- enhanced addressing for session managers [26](#)
- enterprise system connection channel [187](#)
- environment
 - 3270 IDS [301](#), [302](#)
- environment, multiple-network
 - activating resources [483](#)
 - combined APPN and subarea network [414](#)
 - disjointed networks [414](#)
 - network management with NetView [474](#)
 - terminating sessions [491](#)
- ESCON [187](#)
- ESIRFMSG start option [27](#)
- estimated number of active sessions (EAS)
 - APPL definitions [528](#)
 - LFBUF and [542](#)
 - LU definitions [528](#)
 - TSO, specifying for [558](#)
- exchange of ID processing
 - type 2.1 PUs [183](#)
- exit routine
 - session accounting [216](#)
 - session authorization [216](#)
 - warning [216](#)
- explicit route [256](#)
- extended BIND [189](#)
- extended network addressing [464](#)
- extended network addressing (ENA)
 - addressing compatibility [267](#)
 - addressing restrictions [268](#)
 - compatibility, nonextended address [267](#)
 - element address [267](#)
- extended recovery facility (XRF)
 - cryptography and [326](#)
 - description [325](#)
- extended subnetwork boundaries [77](#)
- external communication adapter (XCA)
 - APPN connections [48](#)
 - connection network, defining a connection to [53](#)

- external communication adapter (XCA) (*continued*)
 - tuning statistics [503](#)
 - VCNS and [329](#)
- EZBDVIPA
 - disconnect [381](#)
 - rebuild [381](#)
- EZBEPOR
 - disconnect [384](#)
 - rebuild [383](#)

F

- failure, warning on initiation [323](#)
- filters
 - NetView filter [496](#)
 - session awareness data filter table [495](#)
 - session monitor filter [496](#)
 - VTAM filter [496](#)
- flow control, data [16](#)
- flow control, subarea to network [271](#)
- forced reactivation [493](#)
- FSIRFMSG start option [27](#)

G

- gateway VTAM
 - CDRM definition statement [445](#)
 - defining [444](#)
 - GWPATH definition statement [445](#)
 - NETWORK definition statement [445](#)
- generalized trace facility (GTF) [313](#)
- generic resources
 - and application programs [324](#)
 - and TSO [380](#)
 - configuration, coupling facility failures for [359](#)
 - implementation considerations [358](#)
 - in a sysplex [352](#)
 - maintenance [365](#)
 - mapping [353](#)
 - members, initiating sessions with [357](#)
 - partner LU mapping [354](#)
 - requirements [352](#)
- GTF trace data [313](#)
- GTF trace facility [583](#)
- GWSSCP start option [442](#)

H

- HALT operator command [493](#)
- halting VTAM [493](#)
- Hardware Configuration Definition (HCD) [170](#)
- HAVAIL operand on APPL definition statement [325](#)
- HCD (Hardware Configuration Definition) [170](#)
- high performance data transfer MPC [41](#)
- high performance routing (HPR) [388](#)
- HOST definition statement
 - UNITSZ operand [548](#)
- host physical unit name [22](#)
- host-to-host channel connection, APPN, defining [42](#)
- HOSTPU start option
 - changing [22](#)
 - ISTPUS major node [483](#)
 - multiple-network [442](#)

HOT I/O detection and termination [545](#)
HPDT MPC [41](#)
HPR (high performance routing) [388](#)
HPR start option (APPN) [391](#)
HPRPST start option (APPN) [391](#)

I

I/O (input/output) operations
 analyzing tuning statistics [511](#)
 information gathering [511](#)
IBM Open Systems Adapter
 ATM LAN emulation connections [55](#)
 ATM native connections [55](#)
 LAN connections [55](#)
ICSF (Integrated Cryptographic Service Facility) [593](#)
IDBLK operand on PU definition statement
 type 2.1 PU [183](#)
identifying host, single-domain [21](#)
IDNUM operand on PU definition statement
 type 2.1 PU [183](#)
ILU (independent LU)
 adjacent link station [189](#)
 characteristics [189](#)
 defining [190](#)
 description [189](#)
 dynamic definition of [191](#)
 dynamic selection of sessions [195](#)
 extended BIND [189](#)
 multiple connections [194](#)
 overview [10](#), [189](#)
 restrictions [196](#)
 sessions in a combined network [412](#)
IMS, extended recovery facility (XRF) [325](#)
Incident
 3270 IDS [311](#)
Incident validation [315](#)
independent LU (ILU)
 adjacent link station [189](#)
 characteristics [189](#)
 defining [190](#)
 description [189](#)
 dynamic definition of [191](#)
 dynamic selection of sessions [195](#)
 extended BIND [189](#)
 multiple connections [194](#)
 overview [10](#), [189](#)
 restrictions [196](#)
 sessions in a combined network [412](#)
Information APARs [xxxii](#)
Information Management System, extended recovery facility (XRF) [325](#)
initial status
 overview [483](#)
 warm start and [482](#)
initiation failure [323](#)
initiation, automatic [210](#)
INOPCODE start option [27](#)
INOPDUMP start option [27](#)
input/output (I/O) operations
 analyzing tuning statistics [511](#)
 information gathering [511](#)
installation and definition [581](#)
Integrated Cryptographic Service Facility (ICSF) [593](#)

interchange node (APPN)
 coding [402](#)
 combined APPN and subarea network [402](#)
 overview [7](#)
interconnected networks, sample [450](#)
intermediate routing node [266](#)
Internet, finding z/OS information online [xxxiv](#)
interpret table
 defined [233](#)
 defining for TSO [566](#)
 LOGTAB operand [233](#)
 sample [566](#)
IOBUF buffer pool
 application data transfer [543](#)
 coattailing guidelines [552](#)
 MAXBFRU operand and [548](#)
 overview [537](#)
IOINT start option [27](#)
IOPURGE start option [405](#), [423](#)
ISO standards, NETID [22](#)
ISTAPNPU generic representation [193](#), [386](#)
ISTCDRDY major node [426](#), [483](#)
ISTCOSDF [208](#), [209](#), [260](#)
ISTDEFIN command list [170](#), [172](#), [605](#)
ISTDINFO command list [170](#), [172](#), [601](#)
ISTINCLM [260](#)
ISTMSFLD default message flooding prevention table [487](#)
ISTPDILU major node, initial status and [483](#)
ISTPUS major node
 non-SNA devices and [185](#)
 overview [483](#)
ISTVTCOS [261](#)
ISTVTMEU translation file [609](#), [611](#)
IVTPRM00 [581](#)

J

JCL for the logon manager [586](#)

K

keyboard
 navigation [627](#)
 PF keys [627](#)
 shortcut keys [627](#)

L

LAN (local area network)
 type 2.1 PU and [184](#)
LANG operand [200](#)
LANGTAB operand [200](#)
LBUILD definition statement
 non-SNA devices, defining [185](#)
 sample [185](#)
leased lines, APPN [47](#)
LEN connections, converting to APPN [401](#)
LEN node
 adjacent link station list [413](#)
 combined APPN and subarea network [401](#)
 migrating to APPN nodes [401](#)

- LEN node (*continued*)
 - overview [8](#)
- LFBUF buffer pool
 - application data transfer [543](#)
 - EAS values and [542](#)
 - guidelines for [542](#)
 - overview [537](#)
 - TSO and [542](#)
- license, patent, and copyright information [629](#)
- limited resources, LU 6.2
 - description [335](#)
 - non-LU 6.2 sessions and [336](#)
- limited resources, network management sessions [336](#)
- LIMQSINT operand on APPL definition statement [336](#)
- LIMRES operand on APPL definition statement [336](#)
- line control operand (LNCTL)
 - logon manager [587](#)
- LINE definition statement
 - DELAY operand
 - CTC coattailing and [554](#)
 - inbound coattailing and [547](#)
 - NCP coattailing and [549](#)
 - MODE operand [183](#)
- link stations
 - PRI or SEC indicator [183](#)
 - primary or secondary, determining [183](#)
- link-level role negotiation, type 2.1 PU [183](#)
- LIST option of START command [24](#)
- LISTBKUP (backup start option list) [23](#)
- lists
 - modifying the number of lists [382](#)
- LNCTL operand on the GROUP definition statement
 - logon manager [587](#)
- local area network (LAN)
 - type 2.1 PU and [184](#)
- LOCAL definition statement [185](#)
- locating resources
 - APPN network [13](#)
 - subarea network [12](#)
- logical unit
 - addressing, dependent LU [196](#)
 - dependent
 - addressing [196](#)
 - combined subarea and APPN network [405](#)
 - defining [196](#)
 - overview [9](#)
 - ownership of [196](#)
 - secondary and [197](#)
 - session establishment [197](#)
 - description [183](#)
 - name, unique [466](#)
 - secondary, dependent and [197](#)
 - types of [9](#)
- logon
 - automatic [210](#)
 - passing request [283](#)
 - validating request [283](#)
- LOGON exits, associated LU names [202](#)
- logon manager JCL [586](#)
- logon manager, defining
 - access
 - dependent logical unit [585](#)
 - independent logical unit [585](#)
 - APPL definition statement (operands) [588](#)
- logon manager, defining (*continued*)
 - configuration definition [588](#)
 - defining the logon manager to VTAM [586](#)
 - halting [592](#)
 - installing [585](#)
 - logon manager operation [585](#)
 - logon manager, copying to system [585](#)
 - monitoring [592](#)
 - source statements
 - application password [590](#)
 - buffers, maximum number [591](#)
 - logon procedure [590](#)
 - minimum channel link [590](#)
 - reserved keywords [591](#)
 - subarea configuration [588](#)
 - subareas, maximum number [591](#)
 - TPF applications, number [591](#)
- logon mode table
 - Class of Service and [260](#)
 - combined APPN and subarea network [419](#)
 - DYNMODTB start option [192](#), [204](#), [417](#)
 - entry
 - Class of Service [318](#)
 - DLOGMOD operand [318](#)
 - DYNDLGMD start option [192](#), [204](#)
 - name [467](#)
 - parameters [318](#)
 - session parameters [203](#)
- logon, automatic [210](#)
- LOGTAB operand
 - interpret tables, defined [233](#)
- loop-adapter-attached devices [187](#)
- low-entry network node
 - adjacent link station list [413](#)
 - combined APPN and subarea network [401](#)
 - migrating to APPN nodes [401](#)
 - overview [8](#)
- LPBUF buffer pool
 - CSALIMIT start option [528](#)
 - expansion [528](#)
 - overview [537](#)
- LSIRFMSG start option [27](#)
- LU
 - addressing, dependent LU [196](#)
 - dependent
 - addressing [196](#)
 - combined subarea and APPN network [405](#)
 - defining [196](#)
 - overview [9](#)
 - ownership of [196](#)
 - secondary and [197](#)
 - session establishment [197](#)
 - description [183](#)
 - name, unique [466](#)
 - secondary, dependent and [197](#)
 - types of [9](#)
- LU 6.2
 - APPCCMD macroinstruction, requirement for LU 6.2 [330](#)
 - APPL definition statement
 - APPC operand [330](#)
 - AUTOSES operand [333](#)
 - DDRAINL operand [333](#)
 - DMINWNL operand [333](#)
 - DMINWNR operand [333](#)

LU 6.2 (continued)

APPL definition statement (continued)

- DRESPL operand [333](#)
- DSESLIM operand [333](#)
- SYNCLVL operand [335](#)
- VERIFY operand [334](#)

ATTN exit [335](#)

CNOS [332](#)

conversation [330](#)

description [330](#)

encryption facility [289](#), [334](#)

limited resources [335](#)

LU-LU verification [334](#)

mode name [331](#)

resource verification reduction [331](#)

session limit [332](#)

session-level verification [335](#)

sessions

- automatic activation [333](#)
- characteristics [331](#)
- contention [331](#)
- establishment [330](#)
- negotiating number of [332](#)
- overview [330](#)
- sending data [331](#)
- terminating idle [335](#)
- termination [330](#)
- with LU 6.2 partners [330](#)

sync point services [335](#)

LU definition statement

EAS operand [528](#)

ENCR operand [290](#)

LU names, associated [202](#)

LU-LU verification [334](#)

LUGROUP operand [197](#)

LUSEED operand [197](#)

M

mainframe

- education [xxxii](#)

major node names, list (ATCCONxx) [28](#)

MAXBFRU operand

- coattailing and [548](#)
- guidelines for setting [551](#)

MAXSUBA operand

- extended subarea addressing [268](#)
- gateway NCP [443](#)
- HOSTSA start option and [267](#)
- logon manager JCL [586](#), [589](#)

message flooding prevention [487](#)

message skeleton file [611](#)

messages

- module name, viewing [27](#)
- suppression of [486](#)
- translation of [609](#)
- types [486](#)

messages, USS

- defining [488](#)
- Multicultural support [199](#)
- translation of [609](#)

migration data host (APPN)

- coding [403](#)
- combined APPN and subarea network [403](#)

migration data host (APPN) (continued)

- overview [8](#)

migration, subarea to APPN [401](#)

MMS [609](#)

MODE operand on LINE definition statement [183](#)

MODEENT macroinstruction

- coattailing and [548](#)
- LANG operand [563](#)
- Non-SNA 3270 devices [562](#)
- PSERVIC operand [562](#)
- SNA 3270 devices [562](#)

model application program definitions [278](#)

model CDRSC definition

- coding guidelines [428](#)
- described [428](#)
- resource state requirements [430](#)

model definition [201](#)

model major node [167](#)

model name table

- overview [201](#)
- sample [202](#)

MODETAB operand on APPL definition statement [318](#)

MODIFY CDRM operator command [437](#)

MODIFY CHKPT operator command [242](#)

MODIFY CNOS operator command [333](#)

MODIFY CSM command [583](#)

MODIFY DEFINE operator command, CNOS and [333](#)

MODIFY DR operator command [180](#)

MODIFY TABLE command [491](#)

MODIFY TNSTAT operator command [503](#)

module name, viewing [27](#)

Monitoring CSM [582](#)

monitoring the domain

- DISPLAY command [486](#), [488](#)
- messages
 - suppression of [487](#)
 - types [486](#)

moving PUs and LUs dynamically [174](#)

MPC (multipath channel) connection for hosts

- connection for hosts [40](#)
- tuning I/O [514](#)

MPC, high performance data transfer [41](#)

MPC, HPDT [41](#)

MSGMOD start option [27](#)

MTU [165](#)

Multicultural support

- defining TSO to [569](#)
- description [199](#)
- LANG operand [200](#), [563](#)
- LANGTAB operand [200](#)
- MVS message service [609](#)
- TSO/VTAM user messages [569](#)

multinode persistent sessions

- configuration requirements [368](#)
- configuration, coupling facility failures for [373](#)
- establishing [370](#)
- overview [366](#)
- planned takeover processing [376](#)
- recovery does not occur or complete [379](#)
- using multiple coupling facility structures for [369](#)

multipath channel (MPC) connection for hosts

- connection for hosts [40](#)
- tuning I/O [514](#)

- multipath channel connections, high performance data transfer (HPDT) [41](#)
- multipath channel connections, HPDT [41](#)
- multiple connections for a type 2.1 node [194](#)
- multiple operator consoles [489](#)
- multiple-network environment
 - activating resources [483](#)
 - combined APPN and subarea network [414](#)
 - disjointed networks [414](#)
 - network management with NetView [474](#)
 - terminating sessions [491](#)
- multiple-network, verification [37](#)
- MVS Message Service [609](#)
- MVS operator functions for TSO [569](#)
- MVS system symbols [30](#)

N

- name
 - class-of-service [467](#)
 - logon mode [467](#)
 - subarea-unique, single-domain [30](#)
- name restrictions, single-domain [30](#)
- Native ATM connections
 - connection to the Open Systems Adapter [56](#)
 - port on the Open Systems Adapter [58](#)
 - TGs over SVCs [62](#)
- native operand for subnetwork boundaries [79](#)
- navigation
 - keyboard [627](#)
- NCP
 - APPN connections
 - assigning addresses [386](#)
 - exhausting session control blocks [386](#)
 - maximum number of sessions [386](#)
 - automatic logon (SNA sessions) [209](#)
 - coattailing for [548](#)
 - connection networks [53](#)
 - controlling in CMC environment
 - dynamic switched major node [167](#)
 - CWALL threshold condition [502](#)
 - enterprise system connection channel [187](#)
 - pacing
 - overview [16](#)
 - reducing congestion [16](#)
 - SLODN tuning statistic [502](#)
 - tuning I/O for [504](#), [511](#)
 - tuning statistics [502](#)
- NETID
 - ISO standards for [22](#)
 - single-domain [22](#)
 - start option [22](#), [442](#)
- NETSRVR definition statement, SLUINIT operand [85](#)
- NetView program
 - alias name translation
 - adjacent SSCP tables and [458](#)
 - defining [475](#)
 - overview [466](#)
 - APPL definition statement [474](#)
 - defining to VTAM [474](#)
 - network management with [474](#)
- network accessible units [9](#)
- network identification [443](#)
- network node server (APPN)

- network node server (APPN) (*continued*)
 - list
 - defining [85](#)
 - replacing [86](#)
 - sample [85](#)
 - overview [5](#)
 - resource registration [239](#)
 - supporting SLU-init sessions [85](#)
- network nodes (APPN), overview [4](#)
- network qualified names
 - CDRSCs and [447](#)
 - concurrent sessions with independent LUs [190](#)
 - defining the NetView program and [474](#)
 - overview [459](#)
- Network Registry, SNA [22](#)
- network routing [256](#)
- NMVTLOG start option [27](#)
- NODELST data set [478](#)
- nodes
 - APPN [3](#)
 - subarea [6](#)
 - subarea and APPN [7](#)
 - types of [3](#)
- NODETYPE start option [402](#), [403](#)
- non-SNA devices
 - cluster controllers, defining [185](#)
 - data-stream characteristics [329](#)
 - defining [185](#)
 - loop-adapter-attached [187](#)
- nonnative network connection [184](#), [472](#)
- NQNMODE start option [426](#), [447](#)
- NSRU (unsolicited network services request unit) [495](#)
- NTRI (NCP/Token-Ring Interconnection)
 - defining a connection to a connection network [53](#)

O

- Open Systems Adapter
 - ATM LAN emulation connections [55](#)
 - ATM native connections [55](#)
 - LAN connections [55](#)
- operands, dynamically changing [174](#), [176](#)
- operator commands
 - DISPLAY BFRUSE [500](#)
 - MODIFY TNSTAT [503](#)
- operator-initiated search [238](#)
- OPNDST macroinstruction, OPTCD operand [284](#)
- OPTCD operand on OPNDST macroinstruction [284](#)
- OSA
 - ATM LAN emulation connections [55](#)
 - ATM native connections [55](#)
 - LAN connections [55](#)
- outages, reducing application program [325](#)
- OWNERID [583](#)

P

- pacing control for sessions
 - congestion, route [273](#)
 - defining values, overriding defined pacing counts [222](#)
 - overview [217](#)
- pacing of data flow [16](#)
- pacing, virtual route [270](#)

- parallel sessions [284, 464](#)
- parmlib member, CSM [581](#)
- PARSESS operand on APPL definition statement [284, 464](#)
- PATH definition statement
 - dependent LU requester, defining [406](#)
- path definition, dynamic update
 - description and sample [274](#)
 - restrictions [276](#)
- path update, dynamic [274](#)
- paths
 - forcing a route (APPN) [613](#)
 - logical [258](#)
 - multiple-domain environment [266](#)
 - multiple-network environment [460](#)
 - network routing for subarea nodes [256](#)
 - physical [256](#)
 - routing in a subarea network [256](#)
 - single-domain environment [256](#)
- peak traffic, unusual [16](#)
- performance
 - APPN
 - checkpointing [242](#)
 - improving search performance [238](#)
 - buffer pools specifications [536](#)
 - collecting data [503](#)
 - common service area (CSA) [528](#)
 - input/output (I/O) operations [504, 511](#)
 - introduction [499](#)
 - NCP tuning statistics [502](#)
 - session pacing [217](#)
- performance monitor interface [502](#)
- peripheral link [257](#)
- peripheral nodes
 - attaching to VTAM [184](#)
 - loop-adapter-attached devices [187](#)
 - SNA major node, local [185](#)
- peripheral subnetwork boundaries [77](#)
- PERSIST operand on ACB macroinstruction [327](#)
- persistent LU-LU sessions
 - description [327](#)
 - multinode [366](#)
 - sync point and [335](#)
- physical unit
 - adding dynamically [174](#)
 - attaching to VTAM [184](#)
 - BSBUF buffer pool start option [190](#)
 - deleting dynamically [174](#)
 - loop-adapter-attached [187](#)
 - moving dynamically [174](#)
 - non-SNA devices [185](#)
 - SNA devices
 - defining [185](#)
 - operator-assisted connect [187](#)
 - type 2.1
 - channel, defining [184](#)
 - characteristics [183](#)
 - combined APPN and subarea network [401](#)
 - CPNAME operand and [183](#)
 - identifying [183](#)
 - independent LU [189](#)
 - local area network [184](#)
 - MODE operand and [183](#)
 - nonnative network connection [184](#)
 - restrictions, alias name [190](#)
 - restrictions, independent LU [190](#)
 - session capabilities [189](#)
 - session control block, maintaining [190](#)
 - session establishment [183](#)
 - XID services [183](#)
 - types of [9](#)
- physical unit (*continued*)
 - type 2.1 (*continued*)
 - restrictions, independent LU [190](#)
 - session capabilities [189](#)
 - session control block, maintaining [190](#)
 - session establishment [183](#)
 - XID services [183](#)
 - types of [9](#)
- pools, CSM buffer
 - size of [581](#)
 - source of [581](#)
- PPOLOG start option [27](#)
- predefined cross-domain resource [451](#)
- prerequisite information xxxii
- private storage, allocating [329](#)
- problem diagnosis
 - alerts to NPDA [27](#)
 - module name, viewing [27](#)
 - NMVTLOG start option [27](#)
 - SDLC statistical MDRs [28](#)
 - start options [26](#)
 - time interval notification [27](#)
 - trace facility [28](#)
- program operators
 - associating with USS tables [489](#)
 - defined [494](#)
 - making decisions on setting up [494](#)
- PSERVIC operand [562](#)
- PSSTRACE start option [27](#)
- PU
 - adding dynamically [174](#)
 - attaching to VTAM [184](#)
 - BSBUF buffer pool start option [190](#)
 - deleting dynamically [174](#)
 - loop-adapter-attached [187](#)
 - moving dynamically [174](#)
 - non-SNA devices [185](#)
 - SNA devices
 - defining [185](#)
 - operator-assisted connect [187](#)
 - type 2.1
 - channel, defining [184](#)
 - characteristics [183](#)
 - combined APPN and subarea network [401](#)
 - CPNAME operand and [183](#)
 - identifying [183](#)
 - independent LU [189](#)
 - local area network [184](#)
 - MODE operand and [183](#)
 - nonnative network connection [184](#)
 - restrictions, alias name [190](#)
 - restrictions, independent LU [190](#)
 - session capabilities [189](#)
 - session control block, maintaining [190](#)
 - session establishment [183](#)
 - XID services [183](#)
 - types of [9](#)
- PU definition statement
 - CONNTYPE operand [385](#)
 - CPCP operand [385](#)
 - CPNAME operand, type 2.1 PU and [183](#)
 - IDBLK operand [183](#)
 - IDNUM operand [183](#)
 - PU name and CPNAME uniqueness [401](#)

R

- rapid transport protocol (RTP) [389](#)
- reactivation, forced [493](#)
- reconfiguration, dynamic
 - effect on WARM start [479](#)
 - overview [174](#)
 - VARY ACT, UPDATE technique [176](#)
 - VARY DRDS technique [178](#)
- recovery
 - cryptography and [326](#)
 - description [325](#)
- registering resources [239](#)
- resource control, cross-domain
 - CDRSC definition statement [427](#)
 - controlled by another VTAM domain [427](#)
 - cross-domain resource major node [427](#)
 - dynamic adjacent SSCP table [432](#)
 - locating cross-domain resources [431](#)
 - network resources for XRF [428](#)
 - USERVAR with XRF [428](#)
- resource control, single-domain
 - activating
 - order of activation [482](#)
 - deactivating
 - automatic deactivation of cross-subarea links [492](#)
 - forced deactivation [492](#)
 - forced reactivation [493](#)
 - immediate deactivation [492](#)
 - normal deactivation [492](#)
 - order of deactivation [491](#)
- resource definition statements
 - coding [29](#)
 - sample [30](#)
 - storing [30](#)
 - updating [30](#)
- resource definition, dynamic (model name) [201](#)
- resource registration [239](#)
- resource verification reduction [331](#)
- resource, restoring
 - defined [477](#)
 - MODELST data set [478](#)
 - WARM start [478](#)
- resources in other networks
 - adjacent SSCP (ADJSSCP) table, defining [454](#)
 - adjacent SSCP lists for CDRSCs [431](#)
 - CDRM ownership
 - changing by operator [438](#)
 - changing dynamically [437](#)
 - verifying [437](#)
 - combined APPN and subarea network [410](#)
 - defining [410](#), [425](#)
 - definition statement, VFYOWNER operand [437](#)
 - description [454](#)
 - dynamic [452](#)
 - locating
 - adjacent SSCP table [431](#)
 - ADJCDRM definition statement [432](#)
 - CDRSC routed to each SSCP [431](#)
 - cross-network [454](#)
 - multiple-network configuration [455](#)
 - performance improving [435](#)
 - predefined [450](#)
 - SSCPDYN start option [435](#)
- resources in other networks (*continued*)
 - SSCPORD start option [435](#)
- response time
 - channel-to-channel connection [554](#)
 - coattailing and [527](#), [546](#)
 - DELAY operand and [548](#), [550](#)
 - overview [499](#)
- restoring resource
 - defined [477](#)
 - MODELST data set [478](#)
 - WARM start [478](#)
- restricted materials [621](#)
- restrictions
 - addressing [268](#)
 - independent LU [190](#)
 - name, single-domain [30](#)
 - subarea addressing [270](#)
- RFC (request for comments)
 - accessing online [xxxiv](#)
- route
 - assigning session traffic to a [260](#)
 - explicit [256](#)
 - extension [259](#)
 - planning [264](#)
 - virtual [258](#)
- route congestion [273](#)
- routing
 - APPN
 - broadcast search [236](#)
 - directed search [236](#)
 - directory services [235](#)
 - failure message, CNN routing [253](#)
 - forcing a route [613](#)
 - improving efficiency [238](#), [331](#)
 - local topology database [14](#)
 - network topology database [14](#)
 - topology database [235](#)
 - NETDA/2 (route planning tool) [264](#)
 - network [256](#)
 - overview [13](#)
 - route calculation (APPN), overview [16](#)
 - routing tables [257](#)
 - selection of route, subarea [13](#)
 - topology database (APPN), overview [14](#)
- routing subarea node
 - address incompatibility [268](#)
 - addressing methods [267](#)
 - extended addressing compatibility [270](#)
 - extended network addressing [267](#)
 - path definition, subarea [266](#)
 - subarea addressing restrictions [270](#)
- routing, cross-network
 - communication between network [459](#)
 - COS tables [463](#)
 - extended subarea address gateway [459](#)
 - extended subarea addressing [459](#)
 - gateway NCP [459](#)
 - gateway VTAM [459](#)
 - virtual route, default [463](#)
- routing, network, subarea nodes
 - addressing compatibility [267](#)
 - addressing restrictions [268](#)
 - compatibility, nonextended address [267](#)
 - element address [267](#)

RSIRFMSG start option [28](#)
RTP (rapid transport protocol) [389](#)

S

samples

display buffer use command output [500](#)

networks

multiple-domain [421](#)

multiple-network [440](#)

single-domain [20](#)

tuning statistics report

channel-to-channel [513](#)

SNA controller [511](#)

SATOAPPN COS mapping table [420](#)

SAW (session awareness) data, filtering [495](#)

SDLC statistical MDRs [28](#)

SDLCMDRS start option [28](#)

search reduction [24](#)

searches

broadcast [236](#)

controlling

combined APPN and subarea network [413](#)

order of [413](#)

subarea network searches [414](#)

directed [236](#)

eliminating, for unavailable resources [24](#)

improving performance [242](#)

limiting broadcast searches [414](#)

limiting SNI searches [414](#)

limiting subarea searches [414](#)

operator-initiated [238](#)

performance enhancements [238](#), [331](#)

reducing, for unavailable resources [24](#)

time interval, controlling [405](#), [423](#)

types of [236](#)

security facilities

confidential data [299](#)

single-domain [288](#)

security of data

application program [299](#)

CDRM key [594](#)

conversation-level [334](#)

cross-domain cryptography [594](#)

filing secondary logic unit (SLU) keys [593](#)

session-level [334](#)

single-domain encryption [593](#)

selective termination of idle LU 6.2 sessions

description [335](#)

non-LU 6.2 sessions and [336](#)

selective termination of idle network management sessions [336](#)

session awareness (SAW) data, filtering [495](#)

session initiation request [448](#)

session management exit routine (SME)

adjacent SSCP selection [216](#)

alias selection function [466](#)

for multiple-domain sessions [216](#)

session accounting [217](#)

session authorization [216](#)

when called [216](#)

session managers, enhanced addressing [26](#)

session-partner resource [202](#)

sessions

sessions (*continued*)

accounting for cross-domain [217](#)

acquiring [284](#)

adjacent SSCP selection [435](#)

application program, controlling [210](#)

APPN, setting maximum number [386](#)

associated LU table [202](#)

authorize, establish, or terminate [199](#)

authorizing cross-domain [216](#)

automatic logon [210](#)

backup, extended recovery facility (XRF) [325](#)

between logical units [199](#)

changing an LU's automatic logon specification [211](#)

combined APPN and subarea network [412](#)

control blocks

controlling use of (APPN) [386](#)

exhausting (APPN) [386](#)

controlling [216](#)

CP-CP (APPN), overview [11](#)

cross-network session control, gateway SSCP [444](#)

establishing [216](#)

establishing and terminating with commands [490](#)

establishment

dependent LU [197](#)

type 2.1 PU [183](#)

gateway path selection [460](#)

gateway session accounting [217](#)

LOGON exits [202](#)

logon mode table [203](#)

logon, automatic [210](#)

LU-LU, overview [11](#)

model name table [201](#)

number of half-sessions [217](#)

overview [10](#)

pacing of data flow [16](#)

parallel [284](#)

persistent LU-LU [327](#)

protocol [203](#)

security, LU 6.2 session-level LU-LU verification [334](#)

session accounting exit routine [216](#)

session authorization exit routine [216](#)

session management exit routine [216](#)

SLU-init in network node server (APPN) [85](#)

SNA network interconnection [217](#)

SSCP-SSCP [10](#)

tailoring

authorize or establish [199](#)

between logical units [199](#)

routines used [199](#)

tables used [199](#)

terminating in a multiple-network environment [491](#)

termination [491](#)

USERVAR, session failure [323](#)

SFBUF buffer pool

guidelines for [541](#)

overview [537](#)

shadow resource

application programs [19](#)

CDRSCs [427](#)

overview [438](#)

shortcut keys [627](#)

shut down data transmission [273](#)

sift-down coding [30](#)

single-domain name restrictions [30](#)

- single-domain network
 - overview [19, 421](#)
 - sample [20](#)
- SIRFMSG start option [28](#)
- skeleton file, message [611](#)
- SLODN tuning statistic [502](#)
- slowdown condition
 - analyzing [502](#)
- SLUINIT operand on NETSRVR definition statement [85](#)
- SME (session management exit routine)
 - adjacent SSCP selection [216](#)
 - alias selection function [466](#)
 - for multiple-domain sessions [216](#)
 - session accounting [217](#)
 - session authorization [216](#)
 - when called [216](#)
- SMF [315](#)
- SNA application applicability criteria [302](#)
- SNA devices
 - cluster controllers, defining statically [185](#)
 - loop-adapter-attached [187](#)
- SNA network interconnection (SNI) [441](#)
- SNA Network Registry [22](#)
- SNA protocol specifications [623, 625](#)
- SNAPREQ start option [28](#)
- SNI (SNA network interconnection) [441](#)
- SNVC start option [78](#)
- softcopy information [xxxii](#)
- SONLIM start option [543](#)
- SORDER start option [404, 413](#)
- SPBUF buffer pool, overview [537](#)
- SRCHRED start option [25](#)
- SRCLEAR operand on the MODIFY RESOURCE command [25](#)
- SRCOUNT start option [25](#)
- SRTIMER start option [25](#)
- SSCP
 - identifier [21](#)
 - overview [7](#)
 - owning [467](#)
- SSCP takeover
 - combined APPN and subarea network [410](#)
- SSCP, gateway
 - CDRM definition statement [445](#)
 - defining [444](#)
 - GWPATH definition statement [445](#)
 - NETWORK definition statement [445](#)
- SSCPDYN start option [435](#)
- SSCPID start option [21, 442](#)
- SSCPNAME start option [22, 442](#)
- SSCPORD start option [435](#)
- SSEARCH start option [404, 414](#)
- start option list (ATCSTR00) [20, 22, 183](#)
- start option list, backup (LISTBKUP) [23](#)
- start options
 - APPN
 - APPNCOS [79](#)
 - BN [77](#)
 - BNDYN [78](#)
 - BNORD [78](#)
 - CDSERVR [239](#)
 - CDSREFER [239](#)
 - CONNTYPE [401](#)
 - CPCP [387](#)
 - DIRSIZE [240](#)
 - start options (*continued*)
 - APPN (*continued*)
 - DIRTIME [240](#)
 - DYNADJCP [387](#)
 - HPR [391](#)
 - HPRPST [391](#)
 - INITDB [242](#)
 - NODETYPE [402, 403](#)
 - SNVC [78](#)
 - VERIFYCP [334](#)
 - VFYRED [332](#)
 - VFYREDTI [332](#)
 - VRTG [82](#)
 - VRTGCPCP [82](#)
 - AUTHLEN [196](#)
 - AUTORTRY [212](#)
 - AUTOTI [212](#)
 - CMPMIPS [285](#)
 - combined APPN and subarea network [404](#)
 - DYNDLGMD [192, 204](#)
 - DYNLU [387](#)
 - DYNMODTB [192, 204, 417](#)
 - ENHADDR [25](#)
 - gateway VTAM [442](#)
 - IOINT start option [27](#)
 - IOPURGE [405, 423](#)
 - MSGMOD start option
 - alerts to NPDA [27](#)
 - definition [27](#)
 - VTAM trace facility [28](#)
 - multiple-domain
 - performance improvement [422](#)
 - performance problems [422](#)
 - multiple-network
 - GWSSCP [442](#)
 - HOSTPU [442](#)
 - NETID [442](#)
 - SSCPID [442](#)
 - SSCPNAME [442](#)
 - NODETYPE [402, 403](#)
 - NQNMODE [426](#)
 - overriding [23](#)
 - overview [19, 21](#)
 - performance [422](#)
 - prompting for start options [23](#)
 - single-domain
 - creating start options [24](#)
 - host physical unit name [22](#)
 - LISTBKUP [23](#)
 - NETID [22](#)
 - sample [24](#)
 - source of start options [22](#)
 - SORDER start option [404, 413](#)
 - SRCHRED start option [24](#)
 - SRCOUNT start option [24](#)
 - SRTIMER start option [24](#)
 - SSCP identifier [21](#)
 - SSEARCH start option [404, 414](#)
 - SWNORDER [183](#)
 - XNETALS start option [184](#)
 - start-stop devices, communicating with [329](#)
 - start, warm
 - configuration restart [477](#)
 - initial status and [482](#)

- starting
 - application program
 - name of application and [277](#)
 - security [298](#)
 - VTAM
 - configuration list [28](#)
 - creating start options [24](#)
 - problem diagnosis [26](#)
 - prompting for start options [23](#)
 - sample start options, single-domain [24](#)
 - source of start options [22](#)
 - start options [21](#)
- status information, resource
 - defined [477](#)
 - NODELST data set [478](#)
 - WARM start [478](#)
- storage
 - allocating private for application program [329](#)
 - balancing storage needs [499](#)
 - calculating waste [502](#)
 - monitoring buffer use [500](#)
 - shortage of [529](#)
 - VTAM uses [499](#)
- storage constraints, CSM [582](#)
- storage estimates, APPL EAS [579](#)
- storage limits [581](#)
- storing resource definition statements [30](#)
- subarea node address
 - address incompatibility [268](#)
 - addressing methods [267](#)
 - extended addressing compatibility [270](#)
 - extended network addressing [267](#)
 - path definition, subarea [266](#)
 - subarea addressing restrictions [270](#)
- subarea nodes, network routing for
 - dynamic path update [274](#)
 - extended network and subarea nodes [270](#)
 - nodes, nonextended and extended [267](#)
 - PATH definition statements [256](#)
- subarea to network data control [271](#)
- subarea-unique name, single-domain [30](#)
- substituting COS parameters [261](#)
- SUBSTUT operand on COS macroinstruction [261](#)
- summary of changes xxxvii
- SVCs, ATM, , defining TGs over [62](#)
- switched SDLC connection
 - dynamic definition of [167](#)
- switched virtual channels, ATM, defining TGs over [62](#)
- SWNORDER start option [183](#)
- sync point services, LU 6.2
 - description [335](#)
 - persistent LU-LU sessions and [335](#)
- SYNCLVL operand on APPL definition statement [335](#)
- syntax diagram, how to read xxx
- SYS1.VTAMLST, start options [24](#)
- Sysplex Wide Security Associations [380](#)
- sysplex, dynamic definition of VTAM-to-VTAM connections in a [350](#)
- sysplex, functions provided by VTAM in a [341](#)
- Sysplexports [383](#)
- System Management Facility [315](#)
- System Management Facility (SMF) [315](#)
- system symbols, MVS [30](#)

T

- table, replacing [491](#)
- takeover of resources
 - combined APPN and subarea network [410](#)
- tasks
 - (generic resource, removing)
 - steps [360](#)
- TCP/IP
 - online information xxxiv
- Technotes xxxii
- terminals
 - control address space (TCAS) [557](#)
- time sharing
 - buffers [566](#)
 - cataloged procedures [566](#)
- TNSTAT start option [28](#), [503](#)
- token-ring connection
 - APPN, sample 3172 Nways interconnect controller [48](#)
- topology and routing services (TRS) [235](#)
- topology database (APPN)
 - checkpointing [242](#)
 - creating [235](#)
 - improving performance [242](#)
 - local, overview [14](#)
 - network, overview [14](#)
 - overview [14](#)
 - topology exchange [235](#)
- TP (transmission priority)
 - authorized, for LEN connections [196](#)
 - description [258](#)
- TPF (Transaction Processing Facility)
 - access
 - dependent logical unit [585](#)
 - independent logical unit [585](#)
 - APPL definition statement (operands) [588](#)
 - configuration definition [588](#)
 - defining the logon manager to VTAM [586](#)
 - halting [592](#)
 - installing [585](#)
 - logon manager operation [585](#)
 - logon manager, copying to system [585](#)
 - monitoring [592](#)
 - source statements
 - application password [590](#)
 - buffers, maximum number [591](#)
 - logon procedure [590](#)
 - minimum channel link [590](#)
 - reserved keywords [591](#)
 - subarea configuration [588](#)
 - subareas, maximum number [591](#)
 - TPF applications, number [591](#)
- trace facility [28](#)
- tracing CSM [583](#)
- trademark information [632](#)
- Transaction Processing Facility (TPF)
 - access
 - dependent logical unit [585](#)
 - independent logical unit [585](#)
 - APPL definition statement (operands) [588](#)
 - configuration definition [588](#)
 - defining the logon manager to VTAM [586](#)
 - halting [592](#)
 - installing [585](#)

- Transaction Processing Facility (TPF) *(continued)*
 - logon manager operation [585](#)
 - logon manager, copying to system [585](#)
 - monitoring [592](#)
 - source statements
 - application password [590](#)
 - buffers, maximum number [591](#)
 - logon procedure [590](#)
 - minimum channel link [590](#)
 - reserved keywords [591](#)
 - subarea configuration [588](#)
 - subareas, maximum number [591](#)
 - TPF applications, number [591](#)
- translation, alias name
 - adjacent SSCP tables and [458](#)
 - alias selection function (VTAM) [465](#)
 - defining [467](#)
 - LU 6.2 [470](#)
 - NetView [474](#)
 - overview [466](#)
- transmission groups (TG)
 - boundary function-based (APPN) [39](#)
 - parallel
 - overview [40](#)
 - parallel sessions [273](#)
 - VR-based [80](#)
- transmission priority (TP)
 - authorized, for LEN connections [196](#)
 - description [258](#)
- TRL major node [42](#)
- TRLE operand on the PU definition statement [42](#)
- TRS (topology and routing services) [235](#)
- TSO Operator [569](#)
- TSO/E sessions [557](#)
- TSO/VTAM
 - 2741, TWX, or WTTY [565](#)
 - 3270 characteristics [561](#)
 - 3270 large screen considerations [568](#)
 - 3790/3270 configuration [565](#)
 - definition [557](#)
 - estimating active sessions for (EAS) [528](#)
 - exit routines [567](#)
 - implementing [566](#)
 - interpret table, defining [566](#)
 - LFBUF buffer pool and [542](#)
 - messages, user [569](#)
 - multiple-domain network [558](#)
 - performance [567](#)
 - security messages [567](#)
 - session parameters [561](#)
 - single-domain network [558](#)
 - TCAS application, defining [557](#)
 - TCAS program properties, defining [566](#)
 - translation of messages [609](#)
 - translation table [567](#)
- tuning I/O operations
 - analyzing tuning statistics [511](#)
 - information gathering [511](#)
- tuning statistics
 - blocking outbound PIUs, analyzing [527](#)
 - channel-to-channel adapter [512](#)
 - data transfer, analyzing [527](#)
 - gathering [503](#)
 - host processor use, analyzing [526](#)

- tuning statistics *(continued)*
 - I/O buffer size, analyzing [526](#)
 - performance monitor interface [502](#)
 - read attentions (RDATN) [527](#)
 - record format for SNA controller [511](#), [512](#)
 - slowdown, analyzing [502](#)
 - SNA controllers [504](#), [511](#)
- tuning your network
 - APPN
 - checkpointing [242](#)
 - improving search performance [238](#)
 - buffer pools specifications [536](#)
 - collecting data [503](#)
 - common service area (CSA) [528](#)
 - input/output (I/O) operations [504](#), [511](#)
 - introduction [499](#)
 - NCP tuning statistics [502](#)
 - session pacing [217](#)
- type 4 subarea node [87](#)
- type 5 subarea node [87](#)

U

- UNITSZ operand on HOST definition statement [548](#)
- unsolicited network services request unit (NSRU) [495](#)
- UPARM [613](#)
- updating resource definition statements [30](#)
- user interface
 - ISPF [627](#)
 - TSO/E [627](#)
- user variable (USERVAR), classes
 - class
 - automatic [320](#)
 - user-managed [320](#)
 - type
 - dynamic [321](#)
 - static [321](#)
 - volatile [321](#)
- USERVAR (user variable), classes
 - class
 - automatic [320](#)
 - user-managed [320](#)
 - type
 - dynamic [321](#)
 - static [321](#)
 - volatile [321](#)
- USS commands, defined [488](#)
- USS messages
 - defining [488](#)
 - Multicultural support [199](#)
 - translation of [609](#)
- USS table
 - associating with program operators [489](#)
 - changing, modifying [488](#), [491](#)
 - logon and logoff requests from logical units [233](#)
 - Multicultural support [199](#)
 - operation-level [488](#)
 - session-level [233](#)

V

- VCNS
 - overview [329](#)

- VCNS (*continued*)
 - XCA (3172) and [329](#)
- VCNS operand on APPL definition statement [329](#)
- verification, LU-LU [334](#)
- VERIFYCP start option [334](#)
- verifying a multiple-domain network [37](#)
- VFYOWNER operand on CDRSC definition statement [437](#)
- VFYRED start option [332](#)
- VFYREDTI start option [332](#)
- virtual node [51](#)
- virtual route (VR)
 - coattailing and [554](#)
 - description of [258](#)
- virtual route pacing [270](#)
- VN=YES operand on the ADJCP definition statement [54](#)
- VNGROUP operand on the LINE definition statement [53](#)
- VNGROUP operand on the PORT definition statement [54](#)
- VNNAME operand on the LINE definition statement [53](#)
- VNNAME operand on the PORT definition statement [54](#)
- VR (virtual route)
 - coattailing and [554](#)
 - description of [258](#)
- VR-based TG [80](#)
- VRTG operand on the CDRM definition statement [82](#)
- VRTG start option [82](#)
- VRTGCPCP operand on the CDRM definition statement [82](#)
- VRTGCPCP start option [82](#)
- VSAM configuration restart file
 - associating with major nodes [479](#)
 - CONFGDS operand [479](#)
 - CONFGPW operand [479](#)
- VTAM commands
 - 3270 IDS [309](#)
- VTAM Common Network Services
 - overview [329](#)
 - XCA (3172) and [329](#)
- VTAM internal trace [583](#)
- VTAM restricted materials [621](#)
- VTAM terminal I/O coordinator (VTIOC) [557](#)
- VTAM-to-VTAM connections, dynamic definition of [350](#)
- VTAM, defined as
 - CDRM major node [424](#)
 - CDRM minor node [424](#)
 - cross-domain resource manager (CDRM) [424](#)
 - external CDRM [423](#)
 - host CDRM [423](#)
- VTAM, gateway
 - CDRM definition statement [445](#)
 - defining [444](#)
 - GWPATH definition statement [445](#)
 - NETWORK definition statement [445](#)
- VTAM, online information [xxxiv](#)
- VTAMLST file, sample [30](#)
- VTAMSEG major node [483](#)
- VTAMSEG2 major node [483](#)
- VTIOC (VTAM terminal I/O coordinator) [557](#)

W

- warm start
 - configuration restart [477](#)
 - initial status and [482](#)
- warning
 - changing messages [488](#)

- warning (*continued*)
 - exit routine [216](#)
- window size, adaptive pacing [219](#)

X

- XCA (external communication adapter)
 - APPN connections [48](#)
 - connection network, defining a connection to [53](#)
 - tuning statistics [503](#)
 - VCNS and [329](#)
- XCF
 - dynamic definition of VTAM-to-VTAM connections [350](#)
 - setting up the sysplex environment [341](#)
- XDBUF buffer pool, overview [539](#)
- XID operand for an APPN connection [385](#)
- XID processing
 - type 2.1 PUs [183](#)
- XNETALS start option [184](#)
- XRF (extended recovery facility)
 - cryptography and [326](#)
 - description [325](#)

Z

- z/OS Basic Skills Information Center [xxxii](#)
- z/OS, documentation library listing [633](#)



Product Number: 5655-ZOS

SC27-3672-70

